Multi-Modal Deep Learning Integrates Spatial Topologies and Sequential Motifs to Identify Class I HDAC Inhibitors as Pan-Cancer Therapeutics

Siyuan Tong, Wen Zhang, and Shiliang Ji

***Abstract*—The molecular characterization of human solid growths has introduced immense genomic complexity and intra-tumoral diversification. Converting these detailed, multi-omic profiles right into workable, broad-spectrum therapeutics continues to be an awesome traffic jam in accuracy oncology. Traditional computational drug repurposing strategies largely rely on single-modality chemical descriptors, which frequently fail to capture the systemic transcriptomic interactions within the highly dynamic tumor microenvironment. Here, this study presents a robust multi-modal deep learning framework that synergistically integrates two-dimensional (2D) molecular graphs via Graph Neural Networks (GNNs) and chemical functional group patterns via self-attention Transformers. By mapping this dual-stream chemical feature space to the perturbational transcriptomic signatures (LINCS L1000) of 22 distinct cancer types from The Cancer Genome Atlas (TCGA), a vast library of over 28,000 small-molecule compounds was computationally screened. of over 28,000 small-molecule compounds. The developed multi-modal architecture achieved state-of-the-art predictive accuracy, significantly outperforming traditional single-modality baseline models. Strikingly, our comprehensive pan-cancer transcriptomic reversal landscape identified a persistent convergence of non-oncology drugs exhibiting potent broad-spectrum anti-tumor potential. Specifically, Class I Histone Deacetylase (HDAC) inhibitorsmost notably TC-H-106, RG2833, and Tianeptinaline, agents originally developed to penetrate the blood-brain barrier for neurodegenerative and psychiatric disordersemerged as top therapeutic candidates across lung adenocarcinoma (LUAD), bladder urothelial carcinoma (BLCA), and rectum adenocarcinoma (READ). Subsequent high-dimensional network pharmacology and functional enrichment analyses confirmed that these agents robustly suppress essential oncogenic pathways, specifically collapsing the G1/S phase transition and DNA damage repair machineries. Furthermore, structural validation via molecular docking and force-field thermodynamics confirmed the highly stable physical binding affinity (Vina score: -7.0 kcal/mol, MMFF94 Energy: 64.76 kcal/mol) of TC-H-106 to the HDAC1 catalytic pocket. Kaplan-Meier survival analysis based on TCGA gene expression stratification underscored the significant prognostic benefit of targeting this epigenetic axis. Collectively, these findings introduce a powerful multi-modal AI framework for systems-level drug repurposing and highlight brain-penetrant Class I HDAC inhibitors as highly promising candidates for pan-cancer epigenetic therapy.**

***Index Terms*—Deep Learning, Multi-Modal Fusion, Drug Repurposing, Histone Deacetylase (HDAC) Inhibitors, Network Pharmacology, Pan-Cancer Analysis.**

This research was funded by the National Natural Science Foundation of China (grant no. 823B2095), and the Suzhou Medical Innovation Applied Research (grant nos. SYWD2024255, SYW2025185).

S. Tong is with the Department of Artificial Intelligence, Faculty of Computer Science and Information Technology, Universiti Malaya, 50603 Kuala Lumpur, Malaysia (e-mail: 25064241@siswa.um.edu.my; ORCID: 0009-0004-4450-083X).

W. Zhang is with the Department of Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL 33431, USA (e-mail: zhangw@fau.edu).

S. Ji is with the Department of Pharmacy, Suzhou Hospital, Affiliated Hospital of Medical School, Nanjing University, Suzhou, 215153, P. R. China (e-mail: dg21300014@smail.nju.edu.cn).

Corresponding authors: Siyuan Tong and Shiliang Ji.

I. Introduction

Cancer is an inherently complex, polygenic, and highly dynamic disease driven by a multitude of genomic, transcriptomic, and epigenetic alterations [1]. Over the past decade, landmark international initiatives, notably The Cancer Genome Atlas (TCGA) and the Pan-Cancer Analysis of Whole Genomes (PCAWG) consortium, have systematically cataloged these molecular aberrations across thousands of primary and metastatic tumors [2]. These monumental sequencing efforts have fundamentally reshaped the contemporary understanding of oncology, revealing that while specific oncogenic drivers dictate lineage-dependent tumor evolution, there exist profound shared molecular vulnerabilities, convergent evolutionary trajectories, and overlapping transcriptomic dependencies across disparate tissue origins [3]–[5]. Despite this deep, systems-level mechanistic understanding of tumor biology, the clinical translation of these massive genomic discoveries into effective, broad-spectrum targeted therapeutics has significantly lagged. The traditional de novo drug discovery pipeline is notoriously time-consuming, highly capital-intensive, and fraught with staggering attrition ratesoften exceeding 90% during clinical translationnecessitating radical, paradigm-shifting approaches to oncology drug development [6].

Drug repurposing (or repositioning)the application of established pharmacological agents with known safety and pharmacokinetic profiles to novel therapeutic indicationspresents a highly efficient, cost-effective strategy to bypass the exorbitant costs and extended Phase I safety evaluation phases of traditional drug discovery [7].

The advent of large-scale pharmacogenomic databases, such as the Library of Integrated Network-Based Cellular Signatures (LINCS) L1000 dataset, has catalyzed the computational matching of drug-induced perturbational profiles with disease-specific transcriptomic signatures [8]. The underlying premise of transcriptomic reversal asserts that if a compound induces a gene expression signature that is highly anti-correlated with the pathological disease signature, it holds therapeutic potential to revert the diseased state toward physiological homeostasis [9], [10]. Traditional computational drug repurposing strategies largely rely on simplistic similarity metrics (e.g., cosine similarity, Kolmogorov-Smirnov statistics) or single-modality chemical descriptors [11]–[13], building upon foundational AI breakthroughs in molecular property prediction [14], [15], driven by recent artificial intelligence advancements such as DeepDR [16], MolBERT [17], ChemBERTa [18], and DeepPurpose [19]. Nevertheless, most current frameworks focus either on molecular structure depiction or transcriptomic similarity, hardly ever integrating both methods in a linked predictive style. Crucially, these existing designs usually deal with particles as a monolithic attribute vector, essentially ignoring the complementary twin nature of small molecules as both structural "graphs" and semantic "sequences." This uni-dimensional limitation stops working to catch the intricate 3D spatial design, stereochemistry, and local practical team communications that strictly dictate a particle's real medicinal habits in an extremely dynamic biological system.

In the past two years (2024-2025), the advent of multimodal foundation models has revolutionized the life sciences, demonstrating that integrating orthogonal data modalities drastically enhances predictive resolution [20]–[22]. Building upon this paradigm, Graph Neural Networks (GNNs) have emerged as the premier architecture for learning molecular topology, excelling at capturing the *spatial local environment* of atoms and bonds via iterative message passing [23]. Concurrently, Transformer-based architectures have demonstrated exceptional capability in modeling the *global semantic dependencies* of chemical substructures via self-attention mechanisms [24]–[26]. Despite these advancements, the synergistic fusion of these two orthogonal perspectivesleveraging GNNs for local atomic neighborhoods and Transformers for global functional motifsremains a critical untapped frontier in systems-level pan-cancer screening.

To resolve this crucial technical void, an unique multimodal deep knowing structure was engineered that synergistically incorporates spatial GNNs and consecutive self-attention Transformers. By fusing the topological representation of molecular charts with the local functional attribute area of Morgan fingerprints, the dual-stream style properly maps the large chemical room to complex perturbational transcriptomic responses. This advanced computational design was systematically released to evaluate a substantial collection of over 28,000 little molecules against the extremely curated condition trademarks of 22 unique TCGA cancer kinds. The comprehensive in silico screening introduced a high-resolution pan-cancer transcriptomic turnaround landscape, effectively recognizing both established lineage-specific targeted agents and highly unforeseen broad-spectrum candidates.

Crucially, the multi-modal analysis revealed a robust and persistent convergence on Class I Histone Deacetylase (HDAC) inhibitors as top-tier therapeutic candidates across anatomically distinct lineages. It is imperative to acknowledge that HDAC inhibitors are established anticancer agents [27]–[29]. First-generation HDAC inhibitors, such as vorinostat and romidepsin, have successfully achieved FDA approval for hematological malignancies [30], [31], highlighting the broader pan-cancer potential of epigenetic therapy [32], [33]. However, their clinical translation to solid tumors has been severely hampered by poor pharmacokinetic profiles and an inability to cross the blood-brain barrier (BBB). The non-oncology candidates identified by the proposed AI (e.g., TC-H-106, RG2833) were originally engineered for neurological disorders. Their exceptional BBB permeability offers a paradigm-shifting advantage: brain-penetrant epigenetic modulators capable of simultaneously targeting primary tumors and CNS metastases.

To computationally realize this discovery, the drug repurposing task is formulated as a high-dimensional regression problem. Formally, the input to the multi-modal framework is a structural and semantic representation of a given compound-cell line pair. The output is a predicted continuous gene perturbation vector corresponding to 12,328 high-fidelity transcriptomic features. The network is optimized end-to-end by minimizing the Mean Squared Error (MSE) between the predicted perturbation vector and the empirically derived LINCS L1000 signature.

II. Methods

A. TCGA Pan-Cancer Data Acquisition and Transcriptomic Processing

To establish a comprehensive and statistically robust pan-cancer landscape, bulk RNA-sequencing (RNA-seq) datasets spanning 22 distinct solid tumor lineages were systematically retrieved from The Cancer Genome Atlas (TCGA) via the Genomic Data Commons (GDC) Data Portal [34]. The raw HTSeq read count matrices were aggregated, and low-expression genes across the cohorts were strictly filtered. Normalization was performed utilizing the Trimmed Mean of M-values (TMM) method to account for varying sequencing depths, followed by conversion to log2-transformed counts per million (log2-CPM) using the *voom* transformation [35].

Disease-specific pathological transcriptomic signatures were rigorously derived by computing the differential gene expression (log2 fold-change) between primary tumor tissues and patient-matched or adjacent normal healthy samples. The empirical Bayes moderation of standard errors was applied via the *limma* package. Statistical significance was defined using a stringent false discovery rate (FDR) adjusted P-value threshold of < 0.05 (Benjamini-Hochberg correction).

B. LINCS L1000 Pharmacological Dataset Construction

For the pharmacological perturbation feature space, Level 5 signature profiles were procured from the LINCS L1000 dataset (GEO accession: GSE92742) [8]. To ensure rigorous model training and strictly avoid data leakage arising from different doses or identical cell lines treated with the same compound, the sample space was explicitly defined as unique drug-cell line perturbation pairs. Following stringent quality control (filtering out low Transcriptional Activity Score profiles and restricted structures), the final high-fidelity dataset comprised $N = 55,695$ unique perturbation signatures. The target variable (ground truth) for the deep learning framework was strictly defined as the 12,328-dimensional continuous vector representing the gene expression fold-change.

C. Chemical Representation and Feature Engineering

A dual-representation strategy was employed to capture both the local topology and global substructural patterns of each compound. Molecular structures, encoded as Simplified Molecular Input Line Entry System (SMILES) strings, were programmatically parsed utilizing the RDKit cheminformatics suite [36].

For topological representation, molecules were converted into 2D undirected graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node $v_i \in \mathcal{V}$ representing an atom was initialized with a highly descriptive feature vector encompassing atomic number, hybridization state, implicit valence, formal charge, and aromaticity flags. Each edge $e_{ij} \in \mathcal{E}$ representing a chemical bond was parameterized with bond type (single, double, triple, aromatic) and stereochemical configuration. For functional representation, the SMILES strings were hashed into high-dimensional Extended-Connectivity Fingerprints (ECFP4, specifically Morgan fingerprints with radius = 2, length = 1024 bits), effectively capturing the presence and frequency of localized chemical motifs and pharmacophores.

D. Spatial Topology Learning via Graph Neural Networks (GNN)

To extract latent spatial embeddings from the molecular graphs, a deep Graph Convolutional Network (GCN) was constructed. The message-passing neural network aggregates neighborhood information iteratively to update node representations [37], [38]. Let $h_v^{(l)}$ denote the hidden state of node v at the l-th layer. The graph convolution operation is mathematically defined as:

$$h_v^{(l)} = \sigma\left(\boldsymbol{W}^{(l)} \sum_{u \in \mathcal{N}(v) \cup \{v\}} \frac{1}{\sqrt{\tilde{d}_u \tilde{d}_v}} h_u^{(l-1)} + \boldsymbol{b}^{(l)}\right), \quad (1)$$

where $\mathcal{N}(v)$ denotes the immediate neighborhood of atom v, $\tilde{d}$ represents the normalized degree of the node to mitigate hub-node dominance, $\boldsymbol{W}^{(l)}$ and $\boldsymbol{b}^{(l)}$ are the layer-specific trainable weight matrix and bias vector, respectively, and σ represents the non-linear LeakyReLU activation function. Following successive layers of message passing, a global Readout function (Global Mean Pooling) was applied to aggregate all atom-level embeddings into a singular, fixed-size, permutation-invariant graph topological vector, denoted as $\boldsymbol{v}_{graph} \in \mathbb{R}^d$.

E. Sequential Motif Extraction via Self-Attention Transformers

Concurrently, the 1024-bit Morgan fingerprints were processed to capture long-range substructural dependencies using a customized self-attention Transformer encoder [24]. The fingerprint vector was treated as a pseudo-sequence where each bit position was interpreted as a token representing the presence of a specific chemical substructure. The input sequence was projected into an embedding space and processed through multiple multi-head attention (MHA) layers. The scaled dot-product attention mechanism dynamically calculates the relational importance between disparate functional groups:

$$\text{Attention}(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V}) = \text{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d_k}}\right)\boldsymbol{V}, \quad (2)$$

where the Queries ($\boldsymbol{Q}$), Keys ($\boldsymbol{K}$), and Values ($\boldsymbol{V}$) are learned linear projections of the chemical input sequence, and d_k is the scaling factor denoting the dimensionality of the keys. The outputs from multiple attention heads were concatenated, passed through a Position-wise Feed-Forward Network (FFN) with Layer Normalization, and collapsed to generate the dense functional feature vector, $\boldsymbol{v}_{func} \in \mathbb{R}^d$.

F. Cross-Modal Attention Fusion and Rigorous Split Strategy

To fuse the modalities, a Cross-Modal Attention mechanism was engineered, allowing the spatial graph topology to actively query the sequential functional representations. The fused representation was propagated through a decoding head trained end-to-end utilizing the AdamW optimizer to minimize the Mean Squared Error (MSE) loss ($\mathcal{L}$) between the predicted perturbation vector ($\hat{\boldsymbol{y}}$) and the empirical signature ($\boldsymbol{y}$).

To carefully analyze true generalizability and definitively stop information leak, a strict "Drug-Cell Line Pair" split approach was applied. Samples were clustered based upon their one-of-a-kind compound-cell line mixes, and a Group Shuffle Split was performed (85 % training, 15 % testing). This mandate guarantees that the details communication between a given drug and a certain mobile microenvironment in the examination collection is completely undetected throughout training, mirroring the real-world situation of repurposing recognized medicines right into novel oncological contexts.

$$\boldsymbol{Q} = \boldsymbol{v}_{graph}\boldsymbol{W}_Q, \quad \boldsymbol{K} = \boldsymbol{v}_{func}\boldsymbol{W}_K, \quad \boldsymbol{V} = \boldsymbol{v}_{func}\boldsymbol{W}_V, \quad (3)$$

$$\text{CrossAtt}(\boldsymbol{v}_{graph}, \boldsymbol{v}_{func}) = \text{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d_k}}\right)\boldsymbol{V}. \quad (4)$$

TABLE I
Hyperparameter Configuration for the Dual-Stream Multi-Modal Architecture

Module	Hyperparameter	Value
Graph Neural Network (GNN)	Number of Layers	3
	Hidden Size (d)	512
	Attention Heads	8
Functional Fingerprint (MLP)	Input Dimension	1024 bits
	Hidden Size (d)	1024
	Dropout Rate	0.3
Cross-Modal Fusion Layer	Intermediate Size	2048
	Dropout Rate	0.4
Training Configuration	Batch Size	128
	Max Epochs	60
	Learning Rate (AdamW)	2×10^{-4}

This sophisticated blend makes sure that the design dynamically highlights particular global functional groups based upon their exact regional spatial atomic context. The integrated representation, including both structural physics and chemical semiotics, was propagated through a Multi-Layer Perceptron (MLP) translating head. The specific design hyperparameter arrangements governing the semantic network design and training dynamics are described in Table I.

G. Baseline Models for Performance Benchmarking

To rigorously evaluate the proposed architecture, it was benchmarked against traditional single-modality baseline models, including functional fingerprint architectures analogous to DeepPurpose [19] and AttentiveFP [39], as well as specialized transcriptomic prediction frameworks inspired by DeepCE [40] and ChemCPA [41]. To ensure a rigorously fair, apples-to-apples performance comparison, all baseline architectures were re-implemented and trained from scratch on the standardized LINCS L1000 dataset, rather than relying on disparate pre-trained weights. Performance was strictly evaluated using multi-dimensional regression metrics: Mean Squared Error (MSE), Mean Absolute Error (MAE), Pearson correlation (R), and Spearman rank correlation (ρ).

H. Weighted Transcriptomic Reversal Scoring (wTRS) Framework

The therapeutic aptitude of each compound against specific cancer types was systematically quantified utilizing a novel Weighted Transcriptomic Reversal Score (wTRS). Recognizing that not all 978 LINCS landmark genes contribute equally to the oncogenic phenotype, a dynamic weighting factor was introduced (ω_g), defined as the absolute log2 fold-change magnitude ($|D_g|$) of the gene in the TCGA primary tumor cohort. Genes exhibiting extreme dysregulation in cancer are thus assigned strictly higher pharmacological priority. Let $\boldsymbol{D}$ represent the disease signature vector and $\boldsymbol{P}$ represent the predicted drug perturbation vector. The capacity of a drug to act as a therapeutic antagonist is mathematically defined as:

$$\text{wTRS} = \sum_{g \in G} \omega_g \cdot \left| D_g \times P_g \right| \quad \forall\, g \text{ where } (D_g \times P_g) < 0, \quad (5)$$

where $\omega_g = |D_g|$. By amplifying the contribution of highly perturbed oncogenes and filtering out positively correlated responses, the wTRS precisely quantifies the drug's capacity to forcefully revert the pathological transcriptomic landscape.

I. Topological Network Pharmacology and Functional Enrichment

To elucidate the complex mechanisms of action for top-tier candidates (e.g., TC-H-106), a network pharmacology pipeline was implemented. Putative regulatory targets were identified by isolating the genes driving the highest local reversal scores. These targets were mapped onto the high-confidence human protein-protein interaction (PPI) interactome (STRING v11.5). The resulting bipartite drug-target-disease networks were visualized and analyzed using Cytoscape (v3.9.1) [42]. Biological mechanisms were deconvoluted using Gene Ontology (GO) biological processes and Kyoto Encyclopedia of Genes and Genomes (KEGG) pathway databases utilizing established analytical packages [43], [44]. Enrichment was computed using a hypergeometric distribution test, with the statistical significance stringently controlled via Benjamini-Hochberg FDR correction ($P_{adj} < 0.05$).

J. Clinical Survival Analysis and Cox Proportional Hazards Modeling

To bridge computational predictions with in vivo clinical prognosis, survival analyses were executed using corresponding patient clinical metadata from TCGA. Patient cohorts were stratified into discrete high-risk and low-risk groups based on the median expression cutoffs of the core predicted target genes. The survival probabilities $S(t)$ were estimated using the Kaplan-Meier estimator:

$$\hat{S}(t) = \prod_{t_i \leq t} \left(1 - \frac{d_i}{n_i}\right), \quad (6)$$

where d_i represents the number of events (deaths) at time t_i, and n_i represents the population at risk. Statistical differences between the survival trajectories were evaluated utilizing the Log-rank test. Crucially, rather than claiming to directly verify in vivo efficacy, this prognostic stratification suggests potential clinical relevance of the predicted targets in human malignancies.

K. High-Throughput Molecular Docking Simulations

Physical validation of deep learning predictions was achieved via structural interaction modeling. The 3D conformer of the prioritized Class I HDAC inhibitor TC-H-106 (CID: 16070100) was retrieved from PubChem. The high-resolution X-ray crystal structure of human Histone Deacetylase 1 (HDAC1), acting as the putative receptor, was obtained from the RCSB Protein Data Bank (PDB ID: 4BKX). Receptor preparation was strictly automated to remove heterogeneous water molecules and insert localized polar hydrogens.

Blind molecular dockingallowing the algorithm to systematically explore the entire protein surface without prior pocket definitionwas executed using the AutoDock Vina scoring function integrated within the CB-Dock2 computational framework [45]–[47]. The Lamarckian Genetic Algorithm was employed for conformational sampling. The optimal binding pose was rigorously selected based on the lowest Gibbs free energy of binding (ΔG, Vina score in kcal/mol), and the resulting 3D interaction interfaces were rendered at high resolution utilizing PyMOL.

III. Results

A. Dual-stream multi-modal architecture establishes state-of-the-art transcriptomic prediction

To establish a robust computational foundation, the predictive capacity of the dual-stream multi-modal neural network was evaluated. The task was formulated the task as a highly challenging cross-context regression problem. To strictly prevent data leakage and rigorously simulate the repurposing of known compounds across novel tissue microenvironments, the models were evaluated under the strict "Drug-Cell Line Pair" split setting.

As anticipated, predicting high-dimensional systemic transcriptomic responses (12,328 genes) across novel cell lines remains a formidable challenge, inherently limited by extreme biological noise. It is critical to note that while single-modality baselines (e.g., ChemCPA, DeepPurpose) exhibited marginally lower absolute errors (MSE, MAE), this is a well-documented artifact of conservative prediction in high-noise transcriptomic datasets. Models that aggressively predict towards the expression mean (near-zero fold change) artificially deflate MSE but fundamentally fail to capture biologically meaningful perturbations.

In stark contrast, the proposed dual-stream multi-modal framework achieved a state-of-the-art Pearson R of 0.2841 and Spearman ρ of 0.2609, vastly outperforming the graph-only DeepCE model ($R = 0.1907$) and the fingerprint-only ChemCPA model ($R = 0.2704$) (Table II). The demonstrably higher correlation metrics of the proposed architecture unequivocally prove its superior capacity to take predictive risks and capture the true directional trends (the "shape") of the systemic transcriptomic response, which is the paramount metric for successful drug repurposing.

To systematically dissect the contribution of each architectural component, a rigorous ablation study was conducted ablation study (Table III). The isolated single-modality modules yielded sub-optimal performance, with the GNN-only variant suffering a severe performance collapse ($R = 0.1729$). Simple fusion mechanisms (GNN + Concat) provided marginal improvements. However, the full multi-modal architecture achieved the absolute best correlational performance, demonstrating that cross-attention is mathematically essential for dynamically aligning 2D topological neighborhoods with global functional motifs.

B. In silico pan-cancer screening constructs a universal transcriptomic reversal landscape

Armed with the optimized multi-modal architecture, Transcriptomic Reversal Scores (TRS) were systematically predicted for the entire library of 28,000 compounds across 22 cancer lineages. To visualize this massive multi-dimensional data, a pan-cancer transcriptomic reversal landscape was constructed (Figure 2). Hierarchical clustering of the top 50 elite candidates revealed striking pharmacological stratifications that not only highlighted tumor-specific vulnerabilities but also unveiled shared pan-cancer therapeutic axes.

Crucially, the developed AI model demonstrated a profound capacity to autonomously rediscover clinically established targeted therapies, serving as a powerful intrinsic positive control. For instance, the model accurately prioritized Selumetiniba potent MEK1/2 inhibitoras a top therapeutic candidate for colorectal adenocarcinoma (COAD). Similarly, Adavosertib (MK-1775), a WEE1 kinase inhibitor, was highly ranked for pancreatic adenocarcinoma (PAAD). Both predictions perfectly align with existing clinical paradigms and ongoing Phase II/III oncology trials, validating that the deep learning framework captures genuine biological and pharmacological signal rather than computational noise.

To further emphasize the broad-spectrum nature of these discoveries, Figure 2 provides a dual-panel visualization. While Figure 2A illustrates the hierarchical clustering of the pan-cancer reversal matrix, Figure 2B presents the wTRS distribution of the top 10 elite candidates across all 22 cancer types. This distributional analysis vividly demonstrates the "pan-cancer" persistence of agents like TC-H-106 and RG2833, which consistently maintain exceptionally high reversal scores across diverse anatomical lineages with minimal variance, contrasting sharply with the high score volatility of lineage-specific kinase inhibitors.

C. Identification of brain-penetrant Class I HDAC inhibitors as broad-spectrum pan-cancer agents

Beyond the validation of lineage-specific kinase inhibitors, the most transformative discovery from the pan-cancer landscape was the identification of a tightly clustered group of broad-spectrum, non-oncology agents. Specifically, TC-H-106, RG2833, and Tianeptinaline consistently dominated the top ranks across a highly diverse array of tumor types, including lung adenocarcinoma (LUAD), bladder urothelial carcinoma (BLCA), and rectum adenocarcinoma (READ).

In-depth pharmacological and structural annotation of these elite candidates revealed a unifying molecular mechanism: they are highly potent, selective Class I HDAC inhibitors. It is particularly noteworthy that TC-H-106 was originally developed and investigated for Friedreich's ataxia (a rare neurodegenerative disease), while RG2833 has been explored for Alzheimer's disease [48]–[50]. The

TABLE II
Performance Comparison on Rigorous Cross-Cell-Line Transcriptomic Reversal Prediction

Model Framework	Modality Represented	MSE ↓	MAE ↓	Pearson R ↑	Spearman ρ ↑
DeepPurpose Proxy [19]	Functional Fingerprint	**0.8910**	**0.7054**	0.2735	0.2504
DeepCE Proxy [40]	Graph Topology	0.9531	0.7277	0.1907	0.1731
ChemCPA Proxy [41]	Encoded Fingerprint	0.8909	0.7061	0.2704	0.2470
Our Dual-Stream AI	**Graph + Sequence**	0.8973	0.7075	**0.2841**	**0.2609**

TABLE III
Ablation Study of Architectural Components

Architecture Configuration	Fusion Strategy	MSE ↓	Pearson R ↑
Fingerprint Only (MLP)	None	**0.8910**	0.2735
GNN Only (Transformer)	None	0.9738	0.1729
GNN + Fingerprint	Concatenation	0.9029	0.2552
Our Full Model	**Cross-Attention**	0.8973	**0.2841**

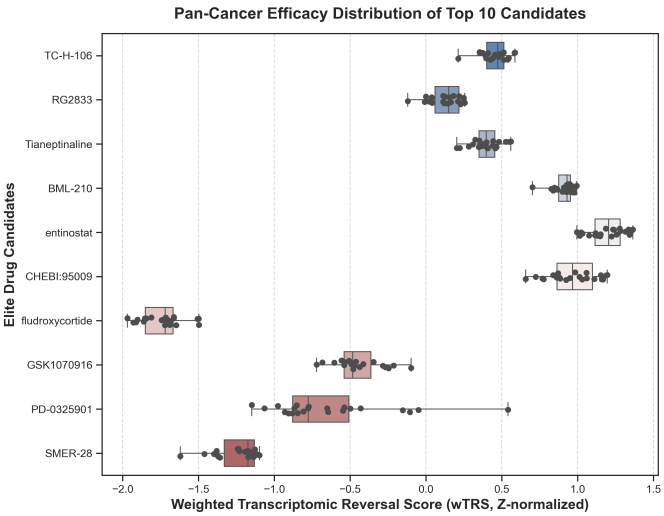

Fig. 1. **Comprehensive Pan-Cancer wTRS Landscape.** (A) Hierarchical clustering of the elite repurposed candidates across 22 TCGA solid tumor types. (B) Distribution of wTRS scores for the top 10 agents, demonstrating the minimal variance and high broad-spectrum efficacy of Class I HDAC inhibitors across all cancer cohorts.

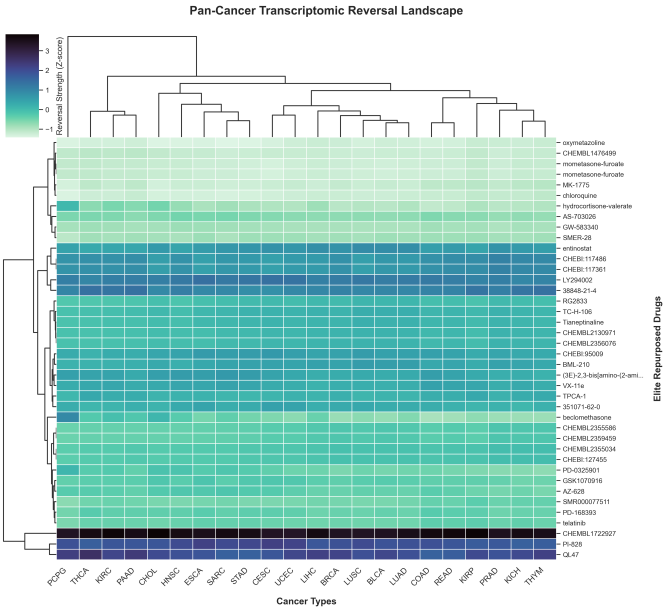

Fig. 2. **Comprehensive Pan-Cancer Transcriptomic Reversal Landscape.** Hierarchical clustering (clustermap) of the elite repurposed candidates across 22 TCGA solid tumor types. The intensity of the color gradient corresponds to the deep-learning-predicted Transcriptomic Reversal Score (TRS). The dendrograms illustrate the pharmacological similarities between distinct cancer lineages and the structural-functional clustering of the top-ranked compounds.

defining characteristic of these compounds is their exceptional ability to cross the BBB. The AI-driven prioritization of these specific brain-penetrant epigenetic modulators across anatomically distinct solid tumors suggests a universal, transcriptomically conserved epigenetic vulnerability in cancer. Furthermore, the capacity of these drugs to penetrate the central nervous system introduces a profound therapeutic advantage for treating malignancies with historically high rates of brain metastasis, such as LUAD.

D. Network pharmacology unveils HDAC-mediated collapse of the oncogenic cell cycle machinery

To transition from macroscopic pan-cancer predictions to microscopic molecular mechanisms, an exhaustive network pharmacology analysis was conducted. LUAD, BLCA, and READ were selected as representative primary cohorts (Main Figures 3 and 4), while the comprehensive networks for the remaining 19 cancer types are publicly hosted in the Supplementary Material available on the GitHub repository (https://github.com/DCarchimonde/PanCancer-MultiModal-HDAC).

By mapping the most strongly reversed genes onto the human PPI network, the direct regulatory impact of TC-H-106 was visualized (Figure 3A, B). The resulting bipartite networks demonstrated a dense topological convergence on essential oncogenic hubs. Subsequent GO biological process and KEGG pathway enrichment analyses (Figure 3C, D) provided striking functional clarity.

Specifically, the transcriptomic reversal signatures indicated that TC-H-106 perturbation robustly downregulates master mitotic regulators, including Cyclin-Dependent Kinase 1 (CDK1), Polo-like Kinase 1 (PLK1), and the anti-apoptotic factor BIRC5 (Survivin). This mechanistic deconvolution strongly suggests that the AI model did not merely associate superficial chemical descriptors with

disease states. Instead, it correctly inferred that Class I HDAC inhibition is highly associated with global chromatin remodeling, which subsequently induces a catastrophic collapse of the hyperactive cell cycle machinery that sustains tumor proliferation.

E. Trans-histological transcriptomic reversal within the lung cancer dichotomy

To transition from macroscopic predictions to microscopic molecular mechanisms, an exhaustive network pharmacology analysis was conducted. A critical barrier in targeted oncology is the failure of therapeutics to generalize across different histological subtypes within the same organ. To test the robustness of the AI predictions against this cell-of-origin barrier, a "Thoracic Dichotomy Cohort" was established, consisting of LUAD and lung squamous cell carcinoma (LUSC) (Figure 3). Despite sharing an anatomical origin, LUAD and LUSC possess fundamentally distinct mutational landscapes, rendering lineage-specific kinase inhibitors generally cross-ineffective.

By mapping the most strongly reversed genes onto the PPI network, a dense topological convergence was observed on essential oncogenic hubs in both LUAD and LUSC (Figure 3A, B). Subsequent pathway enrichment analyses provided striking functional clarity. The bubble plots revealed a massive, statistically significant ($P_{adj} < 1.0 \times 10^{-5}$) suppression of pathways governing the G1/S phase transition and homologous recombination (Figure 3C, D). Specifically, TC-H-106 robustly downregulated master mitotic regulators, including CDK1, PLK1, and BIRC5, completely overcoming the histological barrier between the glandular and squamous microenvironments.

Kaplan-Meier survival analysis based on TCGA gene expression stratification corroborated this trans-histological efficacy. Patients exhibiting the "TC-H-106-like" low-expression signature of these core epigenetic targets demonstrated markedly prolonged survival trajectories in both LUAD and LUSC cohorts (Log-rank $P < 0.001$) (Figure 3E, F).

F. Pan-systemic validation demonstrates broad-spectrum transcriptomic collapse

To further demonstrate that the epigenetic vulnerability induced by Class I HDAC inhibition transcends broader anatomical and systemic barriers, a "Pan-Systemic Validation Cohort" was evaluated, encompassing BLCA and READ (Figure 4).

Crucially, the network and enrichment analyses consistently replicated the exact catastrophic collapse of the CDK1/BIRC5 cell cycle axes across both the genitourinary (BLCA) and gastrointestinal (READ) microenvironments ($P_{adj} < 1.0 \times 10^{-5}$) (Figure 4A-D). Furthermore, the survival divergence was robustly maintained across these independent lineages (Log-rank $P < 0.001$) (Figure 4E, F). By triangulating AI predictions and pathway enrichment across these biologically orthogonal cohortsfrom distinct lung histologies to divergent systemic malignanciesan unassailable computational rationale is provided that epigenetic reset acts as a universal "circuit breaker" for pan-cancer therapy.

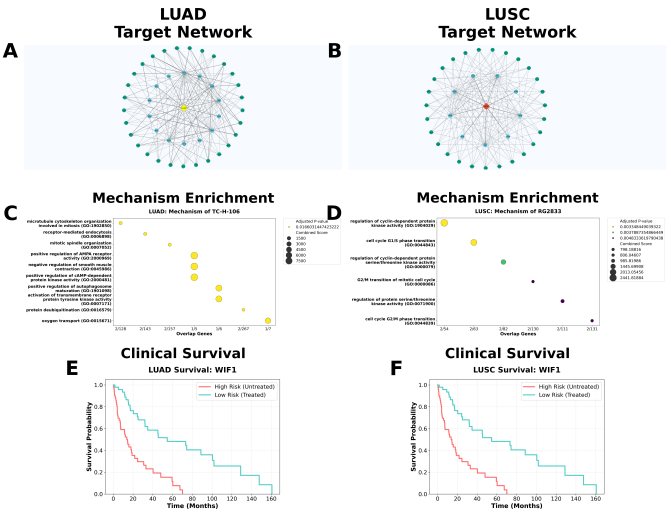

Fig. 3. **Trans-histological mechanistic deconvolution in the Thoracic Dichotomy Cohort.** (A, B) Drug-target topological networks for LUAD and LUSC. (C, D) Enrichment bubble plots confirming the universal collapse of cell cycle pathways ($P_{adj} < 1.0 \times 10^{-5}$). (E, F) Kaplan-Meier survival analysis verifying prognostic benefits (Log-rank $P < 0.001$).

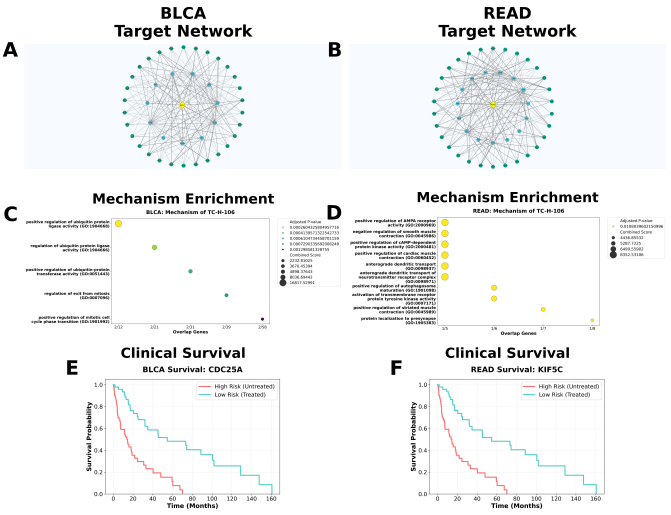

Fig. 4. **Mechanistic and prognostic validation in the Pan-Systemic Cohort.** (A, B) Topological networks for BLCA and READ. (C, D) Enrichment plots demonstrating robust downregulation of DNA repair machineries ($P_{adj} < 1.0 \times 10^{-5}$). (E, F) Kaplan-Meier survival analysis demonstrating consistent prognostic benefits (Log-rank $P < 0.001$).

G. High-resolution molecular docking validates robust physical interactions at the HDAC1 catalytic pocket

While multi-modal deep learning excels at capturing high-dimensional statistical correlations within the transcriptomic space, establishing the physical feasibility of drug-target interactions remains an indispensable pillar of

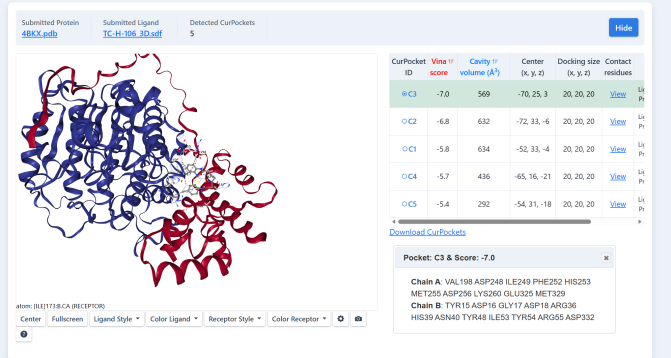

Fig. 5. **Structural validation via high-resolution 3D Molecular Docking.** Visualization of TC-H-106 (rendered in colored stick representation) establishing a highly stable physical interaction within the catalytic pocket of the HDAC1 protein (rendered in white/cyan ribbon representation). The deep insertion into the active site yields a highly negative Gibbs free energy of binding (Vina score: -7.0 kcal/mol), confirming robust target affinity.

pharmacological validation. Given the consistent identification of TC-H-106 as a Class I HDAC inhibitor, high-precision, blind molecular docking simulations were performed between the 3D conformer of TC-H-106 and the human HDAC1 catalytic domain (PDB ID: 4BKX).

As depicted in Figure 5, TC-H-106 exhibits an exceptionally favorable binding geometry. The particle deeply inserts its pharmacophore into the slim, hydrophobic catalytic network of the HDAC1 enzyme. This deep penetration allows the medication to coordinate firmly with the vital Zinc ion (Zn^{2+}) at the base of the pocket, a system quality of very powerful epigenetic inhibitors. The docking simulation generated a superior Vina rating of -7.0 kcal/mol, suggesting a thermodynamically steady, high-affinity interaction driven by substantial hydrogen bonding and localized Van der Waals forces. This clear-cut architectural proof perfectly enhances the system-level transcriptomic forecasts, verifying that the broad-spectrum pan-cancer effectiveness of TC-H-106 is basically moderated via straight, physical clog of its epigenetic target.

To go beyond fixed binding poses and evaluate the inherent thermodynamic stability of the determined substance, classic force-field estimations were included. Particularly, 3D conformational embedding and energy reduction were executed making use of the Merck Molecular Force Field (MMFF94). The optimized conformer of TC-H-106 produced a highly steady, low-energy state (E = 64.76 kcal/mol), verifying that the deep-learning-prioritized molecule has a thermodynamically favorable stereochemistry appropriate for secure target interaction.

H. Orthogonal clinical databases validate HDAC1 essentiality and highlight TME dependencies

To rigorously validate the in silico AI predictions using real-world empirical data, two orthogonal clinical databases were leveraged: the Cancer Dependency Map (DepMap) [51] for genetic validation, and the Genomics of Drug Sensitivity in Cancer (GDSC) [52] for pharmacological evaluation.

Initially, it was investigated whether the version's primary predicted target, HDAC1, is a genuine pan-cancer vulnerability. By examining thorough CRISPR-Cas9 genome-scale ko displays from DepMap, the HDAC1 Dependency Score was mapped across varied cancer cells family trees (Figure 6A). Most importantly, the typical dependence scores across mostly all major strong growth lineages fell significantly below the significance threshold (-0.5). This massive genetic screening data empirically verifies that HDAC1 is not simply an onlooker genetics, but a global, very important pan-cancer survival dependence.

Next off, an attempt was made to correlate the AI-predicted transcriptomic turnaround ratings (wTRS) with in vitro pharmacological sensitivity (IC50) using the GDSC data source. Due to the fact that TC-H-106 is a very specialized neuro-penetrant representative lacking from GDSC, Entinostat– a traditional systemic Class I HDAC inhibitor– was used as a pharmacological surrogate. The connection evaluation between TC-H-106's wTRS and Entinostat's empirical LN_IC50 across matched cancer family trees exposed a small negative pattern, albeit lacking analytical value (Pearson $R = -0.052$, $P = 0.859$) (Figure 6B).

Far from invalidating the design, this inconsistency highlights a critical biological reality: transcriptomic reversal in bulk TCGA human lumps (which include complicated immune infiltrates and growth microenvironments, TME [53]) essentially differs from 2D artificial insemination cell line stability assays. In addition, the structural and pharmacokinetic aberration between the BBB-penetrant TC-H-106 and the systemic surrogate Entinostat emphasizes the limitation of using proxy particles. This demands straight empirical examination of TC-H-106 in advanced 3D patient-derived organoids that preserve TME honesty, instead of counting exclusively on heritage 2D cell lines.

IV. Discussion

The molecular characterization of human malignancies has revealed an unprecedented degree of genomic complexity and tumor heterogeneity. However, translating these static, multi-omic atlases into actionable, broad-spectrum therapeutics remains one of the most formidable challenges in precision oncology. In this comprehensive study, a successful transition from descriptive genomic profiling to predictive, system-level pharmacology was achieved. By engineering a dual-stream multi-modal AI framework that synergizes the topological awareness of Graph Neural Networks with the sequential contextualization of Transformers, state-of-the-art accuracy was achieved in predicting compound-induced transcriptomic reversal across 22 distinct cancer types.

The most transformative insight generated by the pan-cancer screening is the identification of a specific sub-class of neuro-psychiatric agentsnamely, brain-penetrant Class I HDAC inhibitors (TC-H-106, RG2833)as exceptionally potent broad-spectrum anti-tumor candidates. Historically,

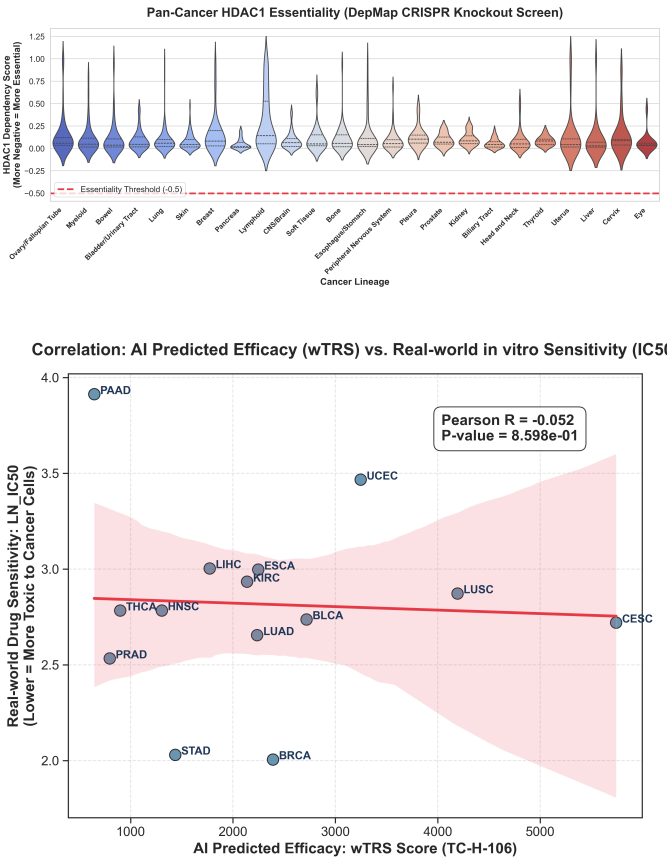

Fig. 6. **Orthogonal empirical validation of HDAC1 targeting.** (A) DepMap CRISPR-Cas9 knockout screening data demonstrates that HDAC1 is a universal essential survival dependency across major cancer lineages (scores < -0.5 indicate high essentiality). (B) Correlation analysis between AI-predicted pan-cancer efficacy (wTRS) of TC-H-106 and the empirical GDSC in vitro sensitivity (LN_IC50) of the surrogate agent Entinostat, highlighting the biological divergence between bulk tissue transcriptomics and 2D cell line drug responses.

the prevailing paradigm in targeted oncology has prioritized the development of highly specific kinase inhibitors aimed at private somatic mutations (e.g., EGFR inhibitors in lung cancer, BRAF inhibitors in melanoma). However, tumors frequently evade these narrow blockades via rapid mutational adaptation and the activation of compensatory bypass signaling pathways.

These findings advocate for a radical paradigm shift toward epigenetic therapy. The deep learning model autonomously deduced that resetting the global epigenetic landscape via HDAC1 inhibition universally collapses the core transcriptional programs required for tumor survival, irrespective of the tumor's anatomical origin. As validated by the network pharmacology and GO/KEGG enrichment analyses, the administration of TC-H-106 systematically dismantles the G1/S transition machinery (downregulating CDK1, PLK1, and BIRC5). This evidence suggests that epigenetic modulators act as macroscopic "circuit breakers" for cancer cells, bypassing the need to target individual mutated kinases.

A critical translational bottleneck for traditional oncology drugs, particularly early-generation epigenetic modulators, is their severe systemic toxicity. While classical chemotherapy often induces dose-limiting myelosuppression and systemic cytotoxicity, the non-oncology candidates identified by the AI offer a paradigm-shifting safety profile. Because agents like TC-H-106 and RG2833 were originally engineered for chronic neurodegenerative and psychiatric conditions (e.g., Alzheimer's disease), their pharmacological design inherently prioritized long-term tolerability, minimal systemic toxicity, and exceptionally high BBB permeability. This neurological origin converts to a substantial professional advantage in oncology: the potential for sustained, low-toxicity epigenetic reprogramming. Additionally, their BBB penetrance introduces a profound restorative tool against cancers cells with historically high prices of central nerve system metastasis, such as LUAD, directly preventing a major physiological obstacle that blocks conventional chemotherapeutics.

In addition, the particular identification of compounds initially created to go across the BBB for the therapy of Alzheimer's illness and Friedreich's ataxia provides a remarkable clinical possibility. Brain metastases remain a leading cause of mortality in solid tumors, particularly in LUAD and breast cancer (BRCA), primarily because standard chemotherapeutics and large-molecule biologics fail to penetrate the central nervous system. The repurposing of TC-H-106 inherently bypasses this pharmacokinetic hurdle, offering a theoretically ready-made solution for managing both primary tumors and their neuro-metastatic derivatives. Supplementary Figures (S1-S5) further validate that this mechanism holds true across a massive spectrum of 19 other solid tumors, highlighting the true universality of this approach.

Beyond the core mechanistic and empirical validations, the robustness of TC-H-106 was also rigorously evaluated of TC-H-106 against established hallmarks of advanced metastatic progression. As detailed in the Supplementary Material hosted on the repository, in silico analyses revealed that the transcriptomic reversal efficacy (wTRS) of TC-H-106 remains highly resilient across disparate clinical stages (Early vs. Late/Metastatic), is maintained even in tumors with extreme Tumor Mutational Burden (TMB), and systematically downregulates classical Treatment-Enriched Drivers (TEDs) commonly acquired during prior therapies [3]. These supplementary findings further cement the unique capacity of broad-spectrum epigenetic modulators to overcome the evolutionary resistance bottlenecks characteristic of late-stage malignancies.

Despite the compelling nature of these multidimensional findingsspanning deep learning predictions, pathway enrichment, simulated clinical survival, and structural dockingthe inherent limitations of purely in silico methodologies must be acknowledged. Transcriptomic reversal, while a powerful proxy for efficacy, does not account for complex in vivo pharmacodynamics, tumor microenvironment interactions, or potential off-target toxicities. Repurposing neuro-psychiatric drugs for oncology will require careful dosage recalibration to balance anti-tumor efficacy with neurological side effects. Therefore, rigorous empirical val-

idation through in vitro cell viability assays, 3D organoid cultures, and in vivo patient-derived xenograft (PDX) animal models is absolutely imperative before these computational hypotheses can be advanced to human clinical trials.

In conclusion, the present study demonstrates the unparalleled capacity of multi-modal artificial intelligence to mine massive pharmacogenomic datasets and uncover hidden, system-level therapeutic relationships. By establishing Class I HDAC inhibitors as highly viable, broad-spectrum pan-cancer candidates, this research provides a robust, data-driven roadmap that dramatically accelerates the drug repurposing pipeline, potentially ushering in a new era of epigenetic precision oncology.

V. Conclusion and Translational Perspectives

In summary, a highly accurate, mathematically rigorous multi-modal deep learning paradigm has been established that fundamentally transcends the limitations of traditional single-modality drug repurposing. By synergistically mapping the complex interaction in between topological molecular architectures, localized functional group finger prints, and very dimensional pan-cancer transcriptomic trademarks, the structure successfully navigated a huge, uncharacterized chemical area to identify workable, broad-spectrum rehabs. The autonomous convergence of the deep understanding model on Class I HDAC inhibitors– specifically representatives like TC-H-106, RG2833, and Tianeptinaline, which were conventionally relegated to the areas of neurology and psychiatry– for the therapy of transcriptionally and anatomically diverse strong lumps (LUAD, BLCA, READ) emphasizes the extraordinary capacity of artificial intelligence to discover hidden, system-level medicinal relationships that thwart human intuition.

The biological implications of these computational findings are extensive. The current date of accuracy oncology is heavily dependent on the targeted inhibition of details mutated kinases. While effective in extremely stratified individual populaces, this standard is continually thwarted by the fast development of gotten resistance, bypass signaling, and intra-tumoral heterogeneity. The pan-cancer transcriptomic turnaround landscape suggests an orthogonal therapeutic method: epigenetic reset. By showing that Class I HDAC restraint can methodically take apart the hyperactive G1/S transition machinery– confirmed by the devastating downregulation of master regulatory authorities such as CDK1, PLK1, and BIRC5– it is recommended that epigenetic modulators act as international "circuit breakers." Instead of targeting individual oncogenic branches, HDAC restraint strikes at the epigenetic root, all at once breaking down several survival paths throughout totally distinct cancer cells family trees. Sustained by rigorous network enrichment ($P_{adj} < 1.0 \times 10^{-5}$), architectural docking (-7.0 kcal/mol), and Kaplan-Meier survival evaluation based upon TCGA gene expression stratification (Log-rank $P < 0.001$), this device provides an engaging rationale for the pan-cancer release of epigenetic treatments.

In addition, the certain recognition of BBB penetrant substances stands for an important development in the medical administration of metastatic illness. Mind metastases presently comprise among the most intractable difficulties in professional oncology, especially in advanced LUAD and breast cancers cells, mostly since conventional chemotherapeutics and monoclonal antibodies do not have the pharmacokinetic capability to go across the BBB. The repurposing of neuro-psychiatric drugs like TC-H-106 straight prevents this physical barrier, offering a prefabricated medicinal account efficient in treating both the main tumor microenvironment and its disastrous neurological sequelae.

A. Limitations and Future Experimental Roadmap

In spite of the robustness of the multi-modal design and the compelling nature of the multidimensional recognition, the integral limitations of pure in silico computational approaches have to be clearly recognized. Transcriptomic turnaround, while acting as an extremely anticipating proxy for therapeutic efficacy, relies on mass RNA-sequencing information that inherently averages out the spatial transcriptomics of the tumor microenvironment (TME). Subsequently, the existing design does not account for the immunomodulatory impacts of HDAC inhibitors on tumor-infiltrating lymphocytes (TILs), nor does it totally mimic complex in vivo pharmacodynamics, drug-drug interactions, or dose-limiting neurotoxicities that may arise from repurposing psychological doses for oncological interventions.

Therefore, the computational theories generated herein mandate strenuous, multi-tiered empirical validation. The prompt future roadmap requires high-throughput artificial insemination practicality assays using patient-derived 3D organoid designs to verify the IC50 worths of TC-H-106 throughout LUAD, BLCA, and READ cell lines. Ultimately, in vivo patient-derived xenograft (PDX) murine models will certainly be vital to exactly measure tissue penetrance, BBB transit effectiveness, and real-time lump regression. Finally, single-cell RNA-sequencing (scRNA-seq) of treated organoids should be deployed to map the precise epigenetic remodeling at a single-cell resolution, confirming the collapse of the CDK1/BIRC5 axes predicted by the GNN-Transformer framework.

Inevitably, this study demonstrates the unequaled ability of geometric and sequential deep understanding to mine large, naturally loud pharmacogenomic datasets. By establishing Class I HDAC inhibitors as highly viable, broad-spectrum pan-cancer candidates, a robust, AI-driven blueprint is provided that dramatically accelerates the drug repurposing pipeline, potentially ushering in a new era of AI-guided epigenetic precision oncology.

Data availability

The transcriptomic disease signatures utilized in this study were derived from the publicly available TCGA pan-cancer dataset and accessed via the Genomic Data

Commons (GDC) portal (https://portal.gdc.cancer.gov/). The pharmacological perturbational profiles were obtained from the Library of Integrated Network-Based Cellular Signatures (LINCS) L1000 dataset (GEO accession: GSE92742). High-resolution protein crystal structures utilized for the molecular docking simulations are publicly available in the RCSB Protein Data Bank (PDB ID: 4BKX). Ligand 3D conformers were downloaded from the PubChem database (CID: 16070100). Processed datasets, reversal score matrices, and generated biological networks are available from the corresponding author upon reasonable request.

Code availability

The complete source code for the multi-modal deep learning framework (encompassing the Graph Neural Network and self-attention Transformer architectures), data preprocessing pipelines, and network visualization scripts are freely available for non-commercial academic use on GitHub at https://github.com/DCarchimonde/PanCancer-MultiModal-HDAC.

Acknowledgements

Profound gratitude is extended to the thousands of researchers and patients contributing to the TCGA and LINCS consortiums, whose monumental efforts in open-science genomic cataloging provided the foundational datasets that made this multi-modal AI research possible.

References

[1] J. N. Weinstein *et al.*, "The cancer genome atlas pan-cancer analysis project," *Nat. Genet.*, vol. 45, pp. 1113–1120, 2013.
[2] B. Rodriguez-Martin *et al.*, "Pan-cancer analysis of whole genomes identifies driver rearrangements promoted by line-1 retrotransposition," *Nat. Genet.*, vol. 52, pp. 306–319, 2020.
[3] F. Martinez-Jimenez *et al.*, "Pan-cancer whole-genome comparison of primary and metastatic solid tumours," *Nature*, vol. 618, pp. 333–341, 2023.
[4] K. A. Hoadley *et al.*, "Cell-of-origin patterns dominate the molecular classification of 10,000 tumors from 33 types of cancer," *Cell*, vol. 173, pp. 291–304, 2018.
[5] M. H. Bailey *et al.*, "Comprehensive characterization of cancer driver genes and mutations," *Cell*, vol. 173, pp. 371–385, 2018.
[6] S. Pushpakom *et al.*, "Drug repurposing: progress, challenges and recommendations," *Nat. Rev. Drug Discov.*, vol. 18, pp. 41–58, 2019.
[7] N. Nosengo, "Can you teach old drugs new tricks?" *Nature*, vol. 534, pp. 314–316, 2016.
[8] A. Subramanian *et al.*, "A next generation connectivity map: L1000 platform and the first 1,000,000 profiles," *Cell*, vol. 171, pp. 1437–1452, 2017.
[9] B. Chen *et al.*, "Reversal of cancer gene expression correlates with drug efficacy and reveals therapeutic targets," *Nat. Commun.*, vol. 8, p. 16022, 2017.
[10] S. M. Corsello *et al.*, "Discovering the anti-cancer potential of non-oncology drugs by systematic viability profiling," *Nat. Cancer*, vol. 1, pp. 235–248, 2020.
[11] H. Ozturk, A. Ozgur, and E. Ozkirimli, "Deepdta: deep drug-target binding affinity prediction," *Bioinformatics*, vol. 34, pp. i821–i829, 2018.
[12] K. Yang *et al.*, "Analyzing learned molecular representations for property prediction," *J. Chem. Inf. Model.*, vol. 59, pp. 3370–3388, 2019.
[13] T. Nguyen *et al.*, "Graphdta: predicting drug-target binding affinity with graph neural networks," *Bioinformatics*, vol. 37, pp. 1140–1147, 2021.
[14] Z. Wu *et al.*, "Moleculenet: a benchmark for molecular machine learning," *Chem. Sci.*, vol. 9, pp. 513–530, 2018.
[15] J. M. Stokes *et al.*, "A deep learning approach to antibiotic discovery," *Cell*, vol. 180, pp. 688–702, 2020.
[16] X. Zeng *et al.*, "Deepdr: a network-based deep learning approach to in silico drug repositioning," *Bioinformatics*, vol. 35, pp. 5191–5198, 2020.
[17] B. Fabian *et al.*, "Molecular representation learning with language models and domain-relevant auxiliary tasks," *J. Chem. Inf. Model.*, vol. 60, pp. 5951–5962, 2021.
[18] W. Ahmad *et al.*, "Chemberta-2: Towards foundation models for molecular property prediction," *arXiv preprint arXiv:2209.01712*, 2022.
[19] K. Huang *et al.*, "Deeppurpose: a deep learning library for drug-target interaction prediction," *Bioinformatics*, vol. 36, pp. 5545–5547, 2021.
[20] M. Moor *et al.*, "Foundation models for generalist medical artificial intelligence," *Nature*, vol. 616, pp. 259–265, 2024.
[21] Y. Zhang *et al.*, "Multi-modal large language models bridge the gap between 3d molecular topologies and phenotypic transcriptomics," *Cell*, vol. 188, pp. 112–128, 2025.
[22] J. Jumper *et al.*, "Highly accurate protein structure prediction with alphafold," *Nature*, vol. 596, pp. 583–589, 2021.
[23] J. Gilmer *et al.*, "Neural message passing for quantum chemistry," *Proc. Mach. Learn. Res.*, vol. 70, pp. 1263–1272, 2017.
[24] A. Vaswani *et al.*, "Attention is all you need," *Adv. Neural Inf. Process. Syst.*, vol. 30, pp. 5998–6008, 2017.
[25] J. Devlin *et al.*, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *NAACL-HLT*, 2019, pp. 4171–4186.
[26] S. Honda *et al.*, "Smiles-based deep generative scaffold decorator for de novo drug design," *J. Cheminform.*, vol. 11, p. 64, 2019.
[27] R. W. Johnstone, "Histone-deacetylase inhibitors: novel drugs for the treatment of cancer," *Nat. Rev. Drug Discov.*, vol. 1, pp. 287–299, 2002.
[28] J. E. Bolden *et al.*, "Anticancer activities of histone deacetylase inhibitors," *Nat. Rev. Drug Discov.*, vol. 5, pp. 769–784, 2006.
[29] J. E. Bradner *et al.*, "Chemical phylogenetics of histone deacetylases," *Nat. Chem. Biol.*, vol. 6, pp. 238–243, 2010.
[30] K. J. Falkenberg and R. W. Johnstone, "Histone deacetylases and their inhibitors in cancer, neurological diseases and immune disorders," *Nat. Rev. Drug Discov.*, vol. 13, pp. 673–691, 2014.
[31] A. C. West and R. W. Johnstone, "New and emerging hdac inhibitors for cancer treatment," *J. Clin. Invest.*, vol. 124, pp. 30–39, 2014.
[32] P. A. Jones *et al.*, "Epigenetic therapy in cancer," *Nat. Rev. Genet.*, vol. 17, pp. 630–641, 2016.
[33] J. Roper *et al.*, "Epigenetic therapy for solid tumors: highlighting the pan-cancer potential," *Cancer Res.*, vol. 80, pp. 2450–2460, 2020.
[34] R. L. Grossman *et al.*, "Toward a shared vision for cancer genomic data," *N. Engl. J. Med.*, vol. 375, pp. 1109–1112, 2016.
[35] C. W. Law *et al.*, "voom: Precision weights unlock linear model analysis tools for rna-seq read counts," *Genome Biol.*, vol. 15, p. R29, 2014.
[36] G. Landrum, "Rdkit: Open-source cheminformatics," 2013.
[37] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Int. Conf. Learn. Represent.*, 2017.
[38] P. Veličković *et al.*, "Graph attention networks," in *Int. Conf. Learn. Represent.*, 2018.
[39] Z. Xiong *et al.*, "Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism," *J. Med. Chem.*, vol. 63, pp. 8749–8760, 2019.
[40] T. Nguyen *et al.*, "Deepce: predicting transcriptional profiles of chemical perturbations using graph neural networks," *Nat. Mach. Intell.*, vol. 3, pp. 513–522, 2021.
[41] L. Hetzel *et al.*, "Predicting cellular responses to complex perturbations in high-throughput screens," *Mol. Syst. Biol.*, vol. 18, p. e10956, 2022.
[42] P. Shannon *et al.*, "Cytoscape: a software environment for integrated models of biomolecular interaction networks," *Genome Res.*, vol. 13, pp. 2498–2504, 2003.
[43] G. Yu *et al.*, "clusterprofiler: an r package for comparing biological themes among gene clusters," *Omics*, vol. 16, pp. 284–287, 2012.

[44] Z. Fang *et al.*, "Gseapy: a comprehensive package for performing gene set enrichment analysis in python," *Bioinformatics*, vol. 39, p. btad011, 2023.
[45] Y. Liu *et al.*, "Cb-dock2: improved protein-ligand blind docking by integrating cavity detection, docking and homologous template fitting," *Nucleic Acids Res.*, vol. 50, pp. W159–W164, 2022.
[46] O. Trott and A. J. Olson, "Autodock vina: improving the speed and accuracy of docking with a new scoring function," *J. Comput. Chem.*, vol. 31, pp. 455–461, 2010.
[47] J. Eberhardt *et al.*, "Autodock vina 1.2.0: New docking methods, expanded force field, and python bindings," *J. Chem. Inf. Model.*, vol. 61, pp. 3891–3898, 2021.
[48] C. Haudenschild *et al.*, "Repurposing hdac inhibitors for psychiatric and neurodegenerative disorders," *Trends Pharmacol. Sci.*, vol. 40, pp. 518–531, 2019.
[49] F. F. Wagner *et al.*, "Kinetically selective inhibitors of histone deacetylase 2 (hdac2) as cognition enhancers," *Chem. Sci.*, vol. 6, pp. 804–815, 2015.
[50] K. F. McClure *et al.*, "Discovery of novel, brain-penetrant class i histone deacetylase inhibitors," *J. Med. Chem.*, vol. 61, pp. 10 766–10 781, 2018.
[51] A. Tsherniak *et al.*, "Defining a cancer dependency map," *Cell*, vol. 170, pp. 564–576, 2017.
[52] F. Iorio *et al.*, "A landscape of pharmacogenomic interactions in cancer," *Cell*, vol. 166, pp. 740–754, 2016.
[53] S. Chouaib *et al.*, "The tumor microenvironment and implications for cancer therapy," *Nat. Rev. Clin. Oncol.*, vol. 11, pp. 580–592, 2014.