

# Mechanism-Aware Prediction of Tissue-Specific Drug Activity via Multi-Modal Biological Graphs

Anonymous authors

Paper under double-blind review

## Abstract

Predicting how small molecules behave across human tissues is essential for targeted therapy development. While some existing models incorporate tissue identity, they treat it as a label—ignoring the underlying biological mechanisms that differentiate tissues. We present Expresso, a multi-modal architecture that predicts tissue-specific molecular activity by modeling how compounds interact with transcriptomic and pathway-level tissue context. Expresso constructs heterogeneous graphs from GTEx data, linking samples, genes, and pathways to reflect expression profiles and curated biological relationships. These graphs are encoded using a hierarchical GNN and fused with frozen molecular embeddings to produce context-aware predictions. A multi-task pretraining strategy—spanning gene recovery, tissue classification, and pathway-level contrastive learning—guides the model to learn mechanistically grounded representations. On nine tissues, Expresso improves mean AUC by up to 27.9 points over molecule-only baselines. Our results demonstrate that incorporating biological structure—not just tissue labels, yields more accurate and interpretable models for tissue-specific drug behavior.

## 1 Introduction

While machine learning has advanced rapidly in molecular modeling, many state-of-the-art models still evaluate drug compounds in isolation—without considering the biological context in which those molecules operate. In reality, a compound’s effect depends not just on its chemical structure but also on the specific tissue it enters, which genes are expressed there, and which biological pathways are active. These factors fundamentally shape how a drug behaves in the body. As a result, models that ignore biological context often struggle to generalize to new tissues or disease states, limiting their reliability in real-world applications.

Some recent approaches attempt to include tissue information, typically by encoding it as a categorical label or an unstructured text input Mamoshina et al. (2016); Ingraham et al. (2019). While this introduces high-level context, it fails to capture the mechanisms that truly differentiate tissues—such as which cellular programs are active, how genes interact in regulatory pathways, and how these dynamics influence a molecule’s activity. Without modeling these mechanistic foundations, predictions can become brittle and overlook critical aspects of tissue-specific drug behavior.

At the same time, large-scale transcriptomic datasets across human tissues have become increasingly accessible. These datasets don’t merely label tissues—they characterize them, offering detailed, sample-level insight into gene activity and biological pathway dynamics. When properly integrated, this wealth of biological structure can serve as a powerful inductive bias for molecular prediction. Yet, most current architectures lack the capacity to exploit this information to generate biologically meaningful predictions.

We address this gap by introducing Expresso: *Expression and Pathway-driven Representation for Enhancing Sample-context and Specific Organ activity prediction*. Expresso is a biologically informed framework that integrates molecular embeddings with graph-based tissue representations built directly from gene expression and curated pathway networks. Instead of abstracting tissue context as a flat label, Expresso constructs a heterogeneous graph per tissue—where samples, genes, and pathways form a structured system. This graph

encodes both expression magnitude and biological relationships, allowing the model to learn how molecules perturb specific tissues based on their internal molecular logic.

The Expresso architecture utilizes hierarchical tissue encoder that performs message passing over sample-gene-pathway graphs. This enables the model to learn rich tissue embeddings conditioned on both observed expression and known biological priors. During training, Expresso is optimized not only for molecular activity prediction but also for self-supervised tissue modeling—reconstructing masked gene expression, classifying tissue identity, and aligning pathway-level representations. This multi-objective setup strengthens generalization and interpretability of the model.

The contributions of this work are threefold: (1) We introduce Expresso, an architecture that integrates molecular structure with rich, expression-driven biological context by leveraging heterogeneous graphs composed of samples, genes, and pathways specific to each tissue—enabling accurate molecular activity prediction in complex biological environments. (2) We develop a novel graph-based tissue encoder, trained via a multi-objective learning framework, that effectively captures both the functional dynamics and structural organization of tissues by jointly modeling real gene expression data and curated pathway priors. (3) We provide new biological insights by identifying key pathways that mediate compound effects across diverse tissues, facilitating more interpretable and biologically grounded predictions.

By embedding molecular prediction in rich, tissue-specific biological context, Expresso brings the field closer to biologically faithful in-silico pipelines capable of modeling compound behavior with organ-level resolution. This enables more informed prioritization of candidate molecules based on their predicted tissue-specific activity, potentially improving the efficiency and focus of subsequent experimental validation. We make our code and data openly available to the community for further study of the problem: <https://anonymous.4open.science/r/EXPRESSO/README.md>.

## 2 Related Work

Recent molecular representation models such as MolBERT Li & Jiang (2021), MolFormer Wu et al. (2023), and MMELON Suryanarayanan et al. (2024) learn rich embeddings from SMILES, improving property prediction tasks like solubility and toxicity. Graph-based models like ChemBERTa Chithrananda et al. (2020) and GROVER Rong et al. (2020) further leverage structural self-supervision and large corpora but lack biological context, limiting their applicability to tissue-specific activity. KA-GNNs Bresson et al. (2025) employ heterogeneous graph representations and kernel-based attention to model complex relational structures within molecules, enhancing the capture of higher-order chemical dependencies. CLAMP Seidl et al. (2023) incorporates assay text for context via contrastive learning, though it uses unstructured biological descriptions. Domain-specific models like BioBERT Lee et al. (2020) excel in biomedical NLP but are not tailored for molecular reasoning. Pathformer Liu et al. (2024) integrates curated gene-pathway and pathway-pathway networks for gene-level prediction but does not model molecules or tissue-specific expression. Some multimodal approaches, such as DeepCDR Wu et al. (2022) and MultiDCP Liu et al. (2020), combine chemical features with cellular multi-omics to predict drug response in cancer cell lines, yet they are limited to in vitro contexts without tissue specificity. In contrast, Expresso uniquely integrates molecular encoders with structured biological knowledge and real tissue-level gene expression, grounding predictions in physiologically relevant context. This enables more accurate, transferable activity predictions across diverse human tissues.

## 3 Problem Definition

We address the problem of *tissue-specific molecular activity prediction*, where the goal is to determine whether a given molecule is active in a particular human tissue. This is framed as a binary classification task, where each input consists of a pair: a *molecule* and a *target tissue*. The model must predict whether the molecule exhibits biological activity in that specific tissue context.

Formally, each prediction instance is defined by: (1) A *molecule*  $m \in \mathcal{M}$ , represented by its SMILES string. (2) A *target tissue*  $t \in \mathcal{T}$ , from a set of human tissues. The model  $f : \mathcal{M} \times \mathcal{T} \rightarrow \{0, 1\}$  is trained to predict whether  $m$  is active in tissue  $t$ .

To accurately model tissue-specific activity, the representation of each tissue must capture its full biological context. Rather than treating tissue as a categorical variable, we aim to build a learned embedding for each tissue that integrates multiple biological components. In particular, we leverage: (1) *Samples*: Each tissue is represented by multiple samples, where each sample corresponds to a gene expression profile derived from a different human donor. Each sample provides a vector of gene expression values (e.g., TPMs), giving insight into the transcriptional landscape of the tissue. (2) *Genes*: These are the measured units in each sample. Each gene has a quantitative expression level. The expression patterns across samples form the basis of the tissue’s molecular signature. (3) *Pathways*: Biological pathways define sets of genes that participate in shared cellular processes or functions. By incorporating pathway membership information (e.g., a binary gene–pathway matrix), we can group genes into functionally meaningful modules that reflect regulatory and signaling mechanisms.

By uniting these components, we create biologically grounded tissue embeddings essential for accurate, context-aware prediction.

## 4 Model Architecture

We present Expresso, a biologically informed model that predicts molecular activity across human tissues by integrating chemical and transcriptomic data. It combines a molecular encoder with a hierarchical graph-based tissue encoder. Below, we detail the molecular encoder, graph construction and message passing, training losses, and the overall model architecture.

### 4.1 Molecular Encoder

To capture the structural properties of molecules, our model uses a pretrained molecular encoder that maps chemical inputs into continuous vector representations. Specifically, we adopt the compound encoder from CLAMP Seidl et al. (2023), which has been trained on large-scale chemical–assay data using contrastive learning. This encoder processes molecular SMILES strings and generates fixed-dimensional embeddings that capture chemical features, which can be predictive of bioactivity when paired with biological context.

Formally, the encoder defines a mapping  $f : \mathcal{M} \rightarrow \mathbb{R}^d$ , where  $\mathcal{M}$  denotes the space of molecules and  $d$  is the embedding dimension. Given a molecule  $m \in \mathcal{M}$ , the output embedding is  $\mathbf{m} = f(m)$ . These embeddings serve as the molecular input to our downstream fusion and prediction modules.

Although we use CLAMP in our implementation, the architecture is modular and compatible with alternative molecular encoders. This design choice allows flexibility in adapting to different data settings or incorporating domain-specific encoders in future applications.

### 4.2 Tissue Graph Construction

To represent the complex biological landscape underlying tissue-specific molecular activity, we construct a heterogeneous graph that captures interactions across multiple biological levels: samples, genes, and pathways. This graph serves as the foundational structure for downstream tissue representation learning. Each tissue is modeled as an independent graph that integrates biological prior knowledge with tissue-specific measurements derived from gene expression data.

#### 4.2.1 Node Types:

The heterogeneous graph includes three distinct node types: (1) *Sample nodes* ( $\mathcal{V}_s$ ), corresponding to individual biological samples. These nodes represent observed instances of gene expression within a tissue and preserve biological variability across samples. (2) *Gene nodes* ( $\mathcal{V}_g$ ), corresponding to protein-coding genes measured in the transcriptomic profile. Each gene node encodes the role of a specific gene in mediating

molecular function. (3) *Pathway nodes* ( $\mathcal{V}_p$ ), representing curated biological pathways, which serve as biologically curated modules that organize genes into coherent functional units, enabling structured reasoning about cellular processes.

Let  $N$ ,  $G$ , and  $P$  denote the number of samples, genes, and pathways respectively in a given tissue graph. The full node set is thus  $\mathcal{V} = \mathcal{V}_s \cup \mathcal{V}_g \cup \mathcal{V}_p$ .

#### 4.2.2 Edge Types:

Edges in the graph encode functional and structural relationships across the biological hierarchy: (1) *Sample–Gene edges* ( $\mathcal{E}_{s \leftrightarrow g}$ ), an undirected bipartite edge is formed between sample node  $v_s \in \mathcal{V}_s$  and gene node  $v_g \in \mathcal{V}_g$  if the gene  $g$  has nonzero expression in that sample  $s$ . This relation captures the biological observation that samples express specific genes. The corresponding binary edge index is encoded as a sparse adjacency matrix  $A_{sg} \in \{0, 1\}^{N \times G}$ , derived directly from the gene expression tensor  $X \in \mathbb{R}^{N \times G}$  by thresholding nonzero entries. (2) *Gene–Pathway edges* ( $\mathcal{E}_{g \leftrightarrow p}$ ) exist between gene nodes  $v_g \in \mathcal{V}_g$  and pathway nodes  $v_p \in \mathcal{V}_p$ , defined according to a binary gene–pathway membership matrix  $M \in \{0, 1\}^{G \times P}$ , where  $M_{gp} = 1$  indicates that gene  $g$  is known to participate in biological pathway  $p$ . This relation encodes structured biological prior knowledge from curated databases, capturing modular and interpretable groupings of genes that function together. These edges define a bipartite graph that links individual molecular measurements to higher-level biological functions. (3) *Pathway–Pathway edges* ( $\mathcal{E}_{p \leftrightarrow p}$ ) model functional dependencies between biological processes by connecting pathway nodes  $v_{p_1}, v_{p_2} \in \mathcal{V}_p$  via a pathway–pathway adjacency matrix  $A_{pp} \in \{0, 1\}^{P \times P}$ , where  $(A_{pp})_{p_1 p_2} = 1$  reflects known crosstalk or regulatory interactions between pathways  $p_1$  and  $p_2$ . These edges capture relationships such as shared downstream targets, feedback loops, or co-activation, enabling information exchange between functionally related modules and supporting more holistic representations of tissue-specific context.

#### 4.2.3 Node Features:

Each node type is initialized with a biologically meaningful feature vector: (1) *Sample features* are computed from the raw gene expression profile. We first apply a simple transformation and normalization to obtain  $\tilde{X} \in \mathbb{R}^{N \times G}$ . Then, each gene expression value is projected to a higher-dimensional space using a feed-forward network, producing sample-specific gene embeddings. These are aggregated across genes to produce a sample feature vector  $h_s \in \mathbb{R}^d$  per node. (2) *Gene features* are initialized using the mean expression of each gene across samples. This results in a  $G$ -dimensional vector, further transformed using the same feed-forward encoder as above, yielding  $h_g \in \mathbb{R}^d$ . (3) *Pathway features* are derived from structural properties of the pathway–pathway network. We compute graph-theoretic statistics such as node degree and clustering coefficient, along with a spectral embedding of the adjacency matrix  $A_{pp}$  (via Laplacian eigenmaps Belkin & Niyogi (2003)). These features are concatenated and standardized, producing a vector for each pathway, which is then projected into the GNN’s hidden dimension:  $h_p \in \mathbb{R}^d$ .

#### 4.2.4 Biological Motivation:

This graph formulation is biologically grounded and designed to capture the multi-scale organization of cellular function. By modeling *samples* explicitly as nodes, the graph preserves inter-sample variation, allowing the tissue representation to reflect within-tissue heterogeneity in gene expression. The inclusion of *gene* and *pathway* nodes enables mechanistic information flow through curated biological structures, supporting both interpretability and generalization. Furthermore, pathway–pathway edges encode regulatory coordination between biological processes, allowing the model to reason at a systems level. Altogether, the constructed heterogeneous graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  provides a biologically meaningful representation of tissue context, which is then processed by a multi-layer heterogeneous GNN to produce compact, informative tissue embeddings for molecular activity prediction.

### 4.3 Graph Message Passing

To transform the constructed heterogeneous graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  into a compact, informative tissue embedding, we employ a hierarchical graph neural network (GNN) architecture that iteratively refines node representations through multi-relational message passing.

#### 4.3.1 Initial Node Embeddings:

For each node type, initial feature vectors are defined as follows: (1) *Sample nodes*  $v_s \in \mathcal{V}_s$  are assigned features derived by projecting normalized gene expression profiles  $\tilde{X} \in \mathbb{R}^{N \times G}$  through a feed-forward encoder, aggregating per-gene embeddings to form a sample representation  $h_s^{(0)} \in \mathbb{R}^d$ . (2) *Gene nodes*  $v_g \in \mathcal{V}_g$  are initialized using the mean expression across samples, transformed via the same encoder, yielding  $h_g^{(0)} \in \mathbb{R}^d$ . (3) *Pathway nodes*  $v_p \in \mathcal{V}_p$  encode structural graph properties of the pathway-pathway network using concatenated graph-theoretic features and spectral embeddings, linearly projected to obtain  $h_p^{(0)} \in \mathbb{R}^d$ .

#### 4.3.2 Layer-wise Message Passing

At GNN layer  $\ell$ , each node  $v$  updates its embedding by aggregating information from neighbors of different types connected via specific edge relations  $r \in \mathcal{R}$ . Formally, for node  $v$ :

$$h_v^{(\ell+1)} = \text{Norm} \left( h_v^{(\ell)} + \sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{N}_r(v)} \frac{1}{|\mathcal{N}_r(v)|} \cdot \text{Conv}_r^{(\ell)}(h_u^{(\ell)}) \right)$$

Here,  $\text{Conv}_r^{(\ell)}$  denotes a relation-specific graph convolution layer (e.g., SAGEConv Hamilton et al. (2017) or GATv2Conv Brody et al. (2021)) tailored to edge type  $r$ ,  $\mathcal{N}_r(v)$  the set of neighbors of  $v$  under relation  $r$ , and Norm a layer normalization operation. Residual connections ensure stability and mitigate over-smoothing across layers.

This multi-relation propagation mechanism integrates signals bidirectionally between samples and genes, between genes and pathways, and among pathways, thereby capturing local molecular activity, functional grouping, and higher-order regulatory crosstalk.

#### 4.3.3 Tissue-level Embedding Aggregation

After  $L$  such layers, sample node embeddings  $\{h_s^{(L)}\}$  represent context-aware molecular profiles within the tissue. To produce a fixed-length tissue embedding  $\mathbf{z} \in \mathbb{R}^d$ , these sample embeddings are pooled via a multi-query attention mechanism:

$$\mathbf{z} = \text{MultiQueryPooling}(\{h_s^{(L)}\}_{s \in \mathcal{V}_s})$$

This pooling adaptively weights samples’ contributions through learned attention queries, allowing the model to capture both dominant and subtle expression patterns relevant for tissue identity and function. Concretely, we parameterize  $Q$  learnable queries, each of which applies standard multi-head attention (Vaswani et al., 2017) over the set of sample embeddings. The resulting  $Q$  context vectors are averaged to obtain  $\mathbf{z}$ .

Overall, this hierarchical GNN encoding framework enables biologically informed integration of heterogeneous molecular and functional information at multiple scales. By explicitly modeling samples, genes, and pathways with distinct features and relations, and leveraging multi-relational convolutions plus attentive pooling, the model learns rich, interpretable tissue representations that reflect both intra-tissue heterogeneity and systemic biological organization.

#### 4.4 Multi-Task Loss Functions

To ensure that the learned tissue embeddings capture meaningful biological signals and generalize well for downstream molecular activity prediction, we train the model using a combination of complementary loss functions. These losses jointly encourage accurate reconstruction of molecular measurements, tissue classification, and robust representation of pathway-level functional modules.

##### 4.4.1 Gene Expression Reconstruction Loss

The heterogeneous graph explicitly models gene expression through sample–gene edges, enabling self-supervised learning by masking a fraction of gene expression values and tasking the model with their reconstruction. Let  $X \in \mathbb{R}^{N \times G}$  denote the normalized gene expression matrix, and  $\mathcal{M} \subseteq \{1, \dots, N\} \times \{1, \dots, G\}$  denote the set of sample–gene index pairs whose expression values are masked during training, randomly sampled across all samples and genes according to a fixed masking ratio. The model predicts masked expression values  $\hat{X}_{i,j}$  via the inner product of learned sample and gene embeddings:

$$\hat{X}_{i,j} = h_{s_i}^\top h_{g_j}, \quad (i, j) \in \mathcal{M}$$

The reconstruction loss is defined as mean squared error over masked entries:

$$\mathcal{L}_{\text{gene}} = \frac{1}{|\mathcal{M}|} \sum_{(i,j) \in \mathcal{M}} (X_{i,j} - \hat{X}_{i,j})^2$$

This loss encourages the model to learn embeddings that preserve detailed gene expression patterns, facilitating recovery of missing molecular signals.

##### 4.4.2 Tissue Classification Loss

To encourage tissue-specific features in the embedding space, we supervise the model to classify tissue identity from the aggregated tissue embedding  $\mathbf{z}$ . Given  $C$  tissue classes and a ground-truth label  $y \in \{1, \dots, C\}$ , a linear classifier  $f_{\text{cls}}$  predicts logits  $\ell \in \mathbb{R}^C$ :

$$\ell = f_{\text{cls}}(\mathbf{z})$$

The tissue classification loss uses cross-entropy:

$$\mathcal{L}_{\text{cls}} = -\log \frac{\exp(\ell_y)}{\sum_{c=1}^C \exp(\ell_c)}$$

This supervision guides the tissue embedding to capture discriminative features relevant to tissue identity.

##### 4.4.3 Pathway Contrastive Loss

To promote pathway embeddings that reflect known functional relationships and are robust to noise, we apply a contrastive loss on two stochastic embeddings of each pathway node obtained via independent passes through the GNN. For each pathway node  $p$ , let  $\mathbf{z}_p^{(1)}, \mathbf{z}_p^{(2)} \in \mathbb{R}^d$  be its embeddings from these independent stochastic passes. Positive pairs  $(\mathbf{z}_p^{(1)}, \mathbf{z}_p^{(2)})$  are contrasted against negative pairs with other pathway embeddings within the batch.

Using cosine similarity  $\text{sim}(\cdot, \cdot)$ , the loss for pathway  $p$  is:

$$\ell_p = -\log \frac{\exp(\text{sim}(\mathbf{z}_p^{(1)}, \mathbf{z}_p^{(2)})/\tau)}{\sum_{q=1}^P \exp(\text{sim}(\mathbf{z}_p^{(1)}, \mathbf{z}_q^{(2)})/\tau)}$$

where  $\tau > 0$  is a temperature hyperparameter. The total pathway contrastive loss is averaged over all pathways:

$$\mathcal{L}_{\text{pathway}} = \frac{1}{P} \sum_{p=1}^P \ell_p$$

This design leverages inherent stochasticity of the model to generate informative positive pairs, enabling robust and functionally meaningful pathway embeddings without requiring explicit graph augmentations.

#### 4.4.4 Activity Prediction Loss

The primary supervised objective is molecular activity prediction. Given a molecular representation and a tissue embedding, the model predicts a scalar activity score. Let  $\mathbf{m}_i$  denote the molecular embedding obtained from the molecular encoder, and  $\mathbf{t}_i$  the tissue embedding produced by the tissue encoder for the  $i$ -th input. These are linearly projected and concatenated to form a joint representation:

$$\mathbf{z}_i = [\mathbf{W}_m \mathbf{m}_i \parallel \mathbf{W}_t \mathbf{t}_i],$$

where  $\mathbf{W}_m$  and  $\mathbf{W}_t$  are learned projection matrices, and  $[\cdot \parallel \cdot]$  denotes vector concatenation. A classifier then maps  $\mathbf{z}_i$  to a logit  $\hat{y}_i \in \mathbb{R}$ , representing the predicted activity score.

Given the ground-truth binary label  $y_i \in \{0, 1\}$ , we minimize a class-weighted binary cross-entropy loss:

$$\mathcal{L}_{\text{activity}} = -\frac{1}{B} \sum_{i=1}^B [\alpha y_i \log \sigma(\hat{y}_i) + (1 - y_i) \log(1 - \sigma(\hat{y}_i))],$$

where  $B$  is the batch size,  $\sigma(\cdot)$  is the sigmoid function, and  $\alpha = \frac{N_{\text{neg}}}{N_{\text{pos}} + \epsilon}$  is the positive-class weight used to mitigate class imbalance. This formulation balances contributions from active and inactive compounds, improving learning stability across tissues with skewed activity distributions. Overall, this loss directly supervises the model to distinguish active from inactive molecular responses in a tissue-specific context.

#### 4.4.5 Overall Objective

The full training objective is a weighted sum of these losses:

$$\mathcal{L} = \mathcal{L}_{\text{activity}} + \lambda_{\text{gene}} \mathcal{L}_{\text{gene}} + \lambda_{\text{cls}} \mathcal{L}_{\text{cls}} + \lambda_{\text{pathway}} \mathcal{L}_{\text{pathway}},$$

where each  $\lambda$  is a tunable hyperparameter controlling the influence of the corresponding auxiliary task. This multi-task formulation enables the model to leverage both labeled activity data and structured biological information, encouraging tissue embeddings to reflect gene-level variation, tissue identity, and pathway-level organization—all while remaining optimized for the end predictive task.

### 4.5 Model Overview and Training

The full model consists of two main components: (1) a molecular encoder, and (2) a tissue encoder that constructs graph-based tissue embeddings from gene expression data integrated with curated biological knowledge. These two embeddings are concatenated and passed through a classifier to predict molecular activity.

We first pre-train the tissue encoder using self-supervised tasks—gene expression reconstruction, tissue classification, and pathway-level contrastive learning—across all tissues without molecular labels. During supervised training, the molecular encoder is frozen while the tissue encoder is fine-tuned. The combined loss, blending binary cross-entropy with self-supervised objectives, promotes biologically grounded and transferable tissue representations.

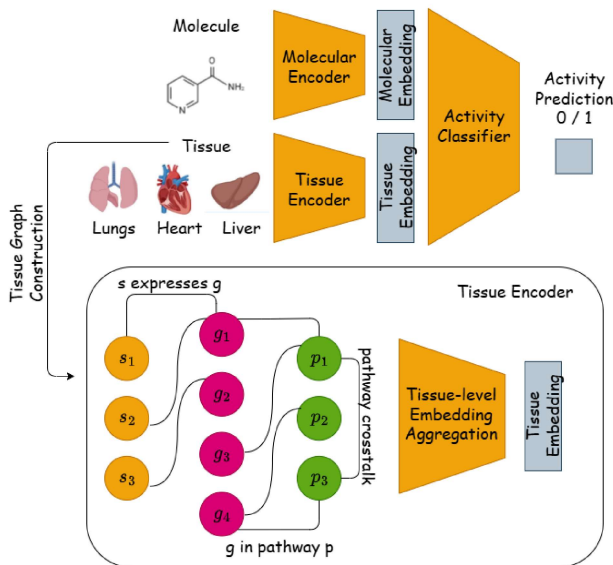


Figure 1: Overview of the Expresso model. A heterogeneous Sample–Gene–Pathway graph captures tissue-specific structure by integrating gene expression, gene–pathway links, and pathway–pathway crosstalk. Node types are color-coded: yellow for samples ( $v_s$ ), pink for genes ( $v_g$ ), and green for pathways ( $v_p$ ).

## 5 Experimental Results

In this section, we empirically evaluate the effectiveness of Expresso for tissue-specific molecular activity prediction. We first present the main results across nine human tissues, followed by an in-depth ablation analysis to dissect the contribution of each model component.

### 5.1 Main Result: Activity Prediction

We benchmark Expresso against six representative baselines spanning language-based, graph-based, and cross-modal paradigms. Full implementation and fine-tuning details for each baseline are provided in Appendix A.3. Additional evaluation metrics and extended results are provided in Appendix B. Tables 1 and 2 summarize the mean AUC scores per tissue. Expresso consistently outperforms all baselines across every tissue, achieving an average AUC of 0.812. This corresponds to a 3.7% relative improvement over the best-performing baseline, BioBERT.

Transformer-based molecular and biochemical language models (BioBERT, MolBERT, MolFormer) consistently achieve strong AUCs in tissues where larger numbers of compound–tissue examples are available (e.g., Liver, Breast), suggesting that these pretrained text- and molecule-based representations capture broadly useful chemical–semantic features when sufficient fine-tuning data are present. Their performance, however, decreases in tissues with limited samples (e.g., Brain, Intestinal) or pronounced class imbalance (e.g., Cervix, Kidney), reflecting the reduced ability of purely sequence- or text-derived encoders to capture tissue-specific biological variation under data-scarce or imbalanced conditions.

Graph-based methods (GROVER, KA-GNN) and the cross-modal baseline (CLAMP) display complementary limitations that explain their comparatively lower performance. Graph encoders effectively capture molecular topology and local substructures but lack mechanisms to integrate tissue- or sample-level transcriptional information; consequently, they struggle to model differential compound activity driven by tissue-specific molecular mechanisms. CLAMP’s multimodal design partially mitigates this by coupling molecular and textual embeddings, yet its contrastive fusion does not fully leverage expression- or pathway-level structure—leading to weaker biological grounding.



Table 1: Performance of Expresso for molecular activity prediction in human tissues compared to various baselines and ablation experiments.

	Model	Brain	Breast	Cervix	Kidney	Liver
	Expresso	<b>0.6471</b> $\pm$ 0.006	<b>0.8233</b> $\pm$ 0.007	<b>0.9325</b> $\pm$ 0.009	<b>0.8764</b> $\pm$ 0.011	<b>0.8207</b> $\pm$ 0.005
Baselines	BioBERT	0.6297 $\pm$ 0.014	0.7867 $\pm$ 0.015	0.8948 $\pm$ 0.008	0.8340 $\pm$ 0.003	0.7796 $\pm$ 0.021
	CLAMP	0.5546 $\pm$ 0.018	0.7849 $\pm$ 0.020	0.8815 $\pm$ 0.007	0.8038 $\pm$ 0.011	0.7712 $\pm$ 0.016
	GROVER	0.5844 $\pm$ 0.012	0.6018 $\pm$ 0.010	0.6060 $\pm$ 0.014	0.6056 $\pm$ 0.009	0.6062 $\pm$ 0.017
	KA-GNN	0.4747 $\pm$ 0.018	0.5292 $\pm$ 0.020	0.5275 $\pm$ 0.007	0.5210 $\pm$ 0.011	0.5477 $\pm$ 0.016
	MolBERT	0.5630 $\pm$ 0.009	0.7760 $\pm$ 0.013	0.8840 $\pm$ 0.011	0.8324 $\pm$ 0.004	0.7777 $\pm$ 0.012
	MolFormer	0.5882 $\pm$ 0.010	0.7886 $\pm$ 0.014	0.8989 $\pm$ 0.012	0.8356 $\pm$ 0.008	0.7834 $\pm$ 0.012
Ablation	MolBert Encoder	0.6387 $\pm$ 0.002	0.7654 $\pm$ 0.006	0.8885 $\pm$ 0.007	0.8323 $\pm$ 0.012	0.7764 $\pm$ 0.013
	ChemGPT Encoder	0.6465 $\pm$ 0.004	0.5699 $\pm$ 0.008	0.6056 $\pm$ 0.006	0.5563 $\pm$ 0.011	0.5762 $\pm$ 0.007
	No Tissue Encoder	0.5756 $\pm$ 0.018	0.7871 $\pm$ 0.014	0.8895 $\pm$ 0.005	0.8061 $\pm$ 0.002	0.7739 $\pm$ 0.017
	Non-Graph Encoder	0.5672 $\pm$ 0.011	0.7855 $\pm$ 0.005	0.8884 $\pm$ 0.002	0.8291 $\pm$ 0.008	0.7801 $\pm$ 0.007
	One-Hot Encoder	0.6076 $\pm$ 0.004	0.8084 $\pm$ 0.006	0.8962 $\pm$ 0.012	0.8534 $\pm$ 0.009	0.7885 $\pm$ 0.004
	Text-Based Encoder	0.6092 $\pm$ 0.018	0.8007 $\pm$ 0.014	0.8993 $\pm$ 0.005	0.8312 $\pm$ 0.002	0.7849 $\pm$ 0.017
	Mean Tissue Expression	0.6176 $\pm$ 0.005	0.7891 $\pm$ 0.009	0.8886 $\pm$ 0.011	0.8325 $\pm$ 0.002	0.7812 $\pm$ 0.013
	No Pretraining	0.6180 $\pm$ 0.010	0.8001 $\pm$ 0.011	0.9110 $\pm$ 0.002	0.8552 $\pm$ 0.009	0.7924 $\pm$ 0.007
	Only $\mathcal{L}_{\text{activity}}$	0.6176 $\pm$ 0.011	0.8097 $\pm$ 0.009	0.8941 $\pm$ 0.010	0.8493 $\pm$ 0.004	0.7938 $\pm$ 0.008
	$\mathcal{L}_{\text{activity}} + \mathcal{L}_{\text{gene}}$	0.6361 $\pm$ 0.014	0.8209 $\pm$ 0.013	0.9106 $\pm$ 0.006	0.8672 $\pm$ 0.015	0.8129 $\pm$ 0.004
	$\mathcal{L}_{\text{activity}} + \mathcal{L}_{\text{cls}}$	0.6134 $\pm$ 0.012	0.8126 $\pm$ 0.015	0.9074 $\pm$ 0.007	0.8481 $\pm$ 0.016	0.7941 $\pm$ 0.007
	$\mathcal{L}_{\text{activity}} + \mathcal{L}_{\text{pathway}}$	0.6218 $\pm$ 0.013	0.8141 $\pm$ 0.006	0.9112 $\pm$ 0.015	0.8553 $\pm$ 0.004	0.8052 $\pm$ 0.007
	Only Sample Nodes	0.5966 $\pm$ 0.006	0.7876 $\pm$ 0.017	0.8979 $\pm$ 0.009	0.8208 $\pm$ 0.015	0.7801 $\pm$ 0.003
	No Gene Nodes	0.6050 $\pm$ 0.003	0.8220 $\pm$ 0.005	0.9054 $\pm$ 0.016	0.8578 $\pm$ 0.013	0.7998 $\pm$ 0.014
	No Pathway Nodes	0.6387 $\pm$ 0.014	0.8091 $\pm$ 0.013	0.9044 $\pm$ 0.005	0.8616 $\pm$ 0.001	0.7949 $\pm$ 0.013

In contrast, Expresso explicitly integrates molecular structure, sample-level gene expression, and pathway topology, enabling context-aware molecular representations that better capture tissue-specific mechanisms of compound activity. This biologically grounded integration yields more accurate and generalizable predictions across diverse human tissues.

## 5.2 Ablation Experiments

To understand the contributions of Expresso’s design components, we conduct an ablation study focused on three key aspects: the encoders architecture, self-supervised objectives, and the biological graph structure.

### 5.2.1 The Role of the Molecular Encoder

This experiment examined how the choice of molecular encoder influences downstream performance. While Expresso employs the CLAMP Seidl et al. (2023) molecular encoder, we replaced it with two alternatives to test different representation strategies: MolBERT Li & Jiang (2021), a transformer trained on SMILES strings, and ChemGPT Frey et al. (2023), a generative model trained on textual molecular data. MolBERT performs reasonably well, showing that SMILES-based pretraining captures relevant chemical semantics, whereas ChemGPT underperforms, suggesting that textual generative models lack sufficient structural precision for accurate molecular activity prediction.

### 5.2.2 The Role of the Tissue Encoder

In this experiment, we tested various tissue representation approaches to understand their impact on model performance. The *One-Hot Encoder* uses simple, learnable one-hot embeddings to represent tissues, while the *Text-Based Encoder* employs frozen CLAMP text embeddings to capture tissue-specific textual information, but without task-specific adaptation. The *Non-Graph Encoder* utilizes precomputed pathway activity profiles from ssGSEA, removing pathway-pathway interactions to capture broad functional trends. Additionally, we evaluated *Mean Tissue Expression*, where each tissue is represented by the average gene expression across all samples in that tissue, capturing only the tissue’s general activity profile.

Table 2: Performance of Expresso for molecular activity prediction in human tissues compared to various baselines and ablation experiments.

	Model	Lung	Ovary	Prostate	Skin
	Expresso	<b>0.8763</b> $\pm 0.008$	<b>0.7466</b> $\pm 0.004$	<b>0.6953</b> $\pm 0.014$	<b>0.8950</b> $\pm 0.011$
Baselines	BioBERT	0.8526 $\pm 0.006$	0.7234 $\pm 0.016$	0.6848 $\pm 0.025$	0.8622 $\pm 0.001$
	CLAMP	0.8411 $\pm 0.018$	0.6899 $\pm 0.007$	0.6547 $\pm 0.004$	0.8648 $\pm 0.020$
	GROVER	0.6017 $\pm 0.013$	0.6020 $\pm 0.006$	0.6208 $\pm 0.015$	0.6017 $\pm 0.005$
	KA-GNN	0.5737 $\pm 0.018$	0.5291 $\pm 0.007$	0.5504 $\pm 0.004$	0.5507 $\pm 0.008$
	MolBERT	0.8582 $\pm 0.010$	0.7137 $\pm 0.008$	0.6753 $\pm 0.017$	0.8672 $\pm 0.014$
	MolFormer	0.8458 $\pm 0.015$	0.7049 $\pm 0.004$	0.6588 $\pm 0.017$	0.8683 $\pm 0.009$
Ablation	MolBert Encoder	0.8534 $\pm 0.003$	0.7032 $\pm 0.005$	0.6635 $\pm 0.009$	0.8741 $\pm 0.006$
	ChemGPT Encoder	0.5247 $\pm 0.015$	0.5812 $\pm 0.004$	0.5854 $\pm 0.002$	0.5866 $\pm 0.012$
	No Tissue Encoder	0.8173 $\pm 0.007$	0.6954 $\pm 0.008$	0.6781 $\pm 0.013$	0.8413 $\pm 0.002$
	Non-Graph Encoder	0.8321 $\pm 0.011$	0.7149 $\pm 0.005$	0.6509 $\pm 0.002$	0.8480 $\pm 0.007$
	One-Hot Encoder	0.8568 $\pm 0.005$	0.7425 $\pm 0.011$	0.6862 $\pm 0.010$	0.8649 $\pm 0.003$
	Text-Based Encoder	0.8499 $\pm 0.002$	0.7222 $\pm 0.005$	0.6994 $\pm 0.007$	0.8718 $\pm 0.012$
	Mean Tissue Expression	0.8416 $\pm 0.006$	0.7100 $\pm 0.001$	0.6718 $\pm 0.014$	0.8666 $\pm 0.009$
	No Pretraining	0.8579 $\pm 0.013$	0.7242 $\pm 0.001$	0.6759 $\pm 0.012$	0.8663 $\pm 0.004$
	Only $\mathcal{L}_{\text{activity}}$	0.8557 $\pm 0.012$	0.7147 $\pm 0.009$	0.6707 $\pm 0.011$	0.8743 $\pm 0.010$
	$\mathcal{L}_{\text{activity}} + \mathcal{L}_{\text{gene}}$	0.8695 $\pm 0.003$	0.7322 $\pm 0.014$	0.6738 $\pm 0.015$	0.8812 $\pm 0.013$
	$\mathcal{L}_{\text{activity}} + \mathcal{L}_{\text{cls}}$	0.8592 $\pm 0.015$	0.7132 $\pm 0.004$	0.6733 $\pm 0.013$	0.8743 $\pm 0.016$
	$\mathcal{L}_{\text{activity}} + \mathcal{L}_{\text{pathway}}$	0.8695 $\pm 0.013$	0.7200 $\pm 0.002$	0.6786 $\pm 0.014$	0.8731 $\pm 0.015$
	Only Sample Nodes	0.8370 $\pm 0.007$	0.7016 $\pm 0.014$	0.6803 $\pm 0.016$	0.8578 $\pm 0.005$
	No Gene Nodes	0.8626 $\pm 0.016$	0.7372 $\pm 0.003$	0.6717 $\pm 0.004$	0.8776 $\pm 0.016$
	No Pathway Nodes	0.8596 $\pm 0.007$	0.7328 $\pm 0.012$	0.6559 $\pm 0.013$	0.8733 $\pm 0.015$

The intuition behind these variants was to evaluate how different levels of tissue context affect performance. The results show that Expresso consistently achieves the highest performance across all tissue, while removing the tissue encoder leads to a marked decline. The *One-Hot* and *Text-Based* variants partially recover this loss but remain below Expresso, showing that simple or frozen representations are insufficient to capture complex biological variability. The *Mean Tissue Expression* baseline performs comparably to these simplified variants, indicating that averaged transcriptomic signals provide only coarse tissue-level cues.

### 5.2.3 The Role of Pretraining

Without pretraining, the tissue encoder lacks the opportunity to learn how to represent tissues in a generalizable way. It is only optimized to support the downstream molecular activity prediction task, which limits its ability to model gene and pathway interactions more broadly. Pretraining provides this missing context by exposing the encoder to diverse tissues and self-supervised objectives, enabling it to build richer, biologically grounded representations independent of activity labels.

### 5.2.4 The Role of Auxiliary Losses

We also examine how different self-supervised objectives affect learning. Removing all auxiliary losses reduces performance across the board, confirming that the model benefits from more than just label supervision. Gene-level masking helps the model capture fine-grained expression dynamics, while sample classification reinforces broader tissue-level structure. Pathway-based objectives are especially useful in tissues where higher-order biological processes dominate, enabling the model to reason more effectively about drug activity.

### 5.2.5 The Role of Graph Structure

The architecture of the tissue graph plays a central role in Expresso’s success. When only sample nodes are used, the model lacks meaningful biological grounding. Removing gene nodes, by dropping all sample–gene and gene–pathway edges, eliminates crucial expression-level detail, while excluding pathway nodes strips away functional hierarchy. The full graph—composed of samples, genes, and pathways—enables both




Figure 2: Heatmap of normalized pathway activity across tissues. Each row corresponds to a pathway grouped by functional category, and each column represents a tissue.

bottom-up signal integration and top-down semantic guidance. Each node type contributes a complementary perspective, and removing any of them weakens the model’s ability to reason about tissue-specific drug effects.

## 6 Uncovering Tissue-Specific Insights Through Pathways and Predictions

In this section, we present deeper insights into how the model captures tissue-specific drug behavior. We analyze the biological pathways it relies on across tissues, and evaluate its ability to prioritize therapeutically relevant compounds, offering both interpretability and predictive validation.

### 6.1 Pathway-Level Insights Across Tissues

To better understand how Expresso makes tissue-specific predictions, we conducted an interventional pathway knockout analysis. For each tissue, we systematically removed individual biological pathways from the model’s internal graph and measured the resulting decline in predictive performance. The magnitude of this decrease reflects each pathway’s relative causal contribution to model accuracy, providing an interpretable measure of functional importance.

Figure 2 summarizes the normalized distribution of top pathway categories across tissues, offering a high-level view of tissue-specific functional emphasis. The model’s prioritization patterns align closely with known biological roles. Metabolic pathways are strongly enriched in the liver and kidney—organs central to detoxification and metabolic homeostasis—with moderate representation in breast tissue due to its high biosynthetic activity. Signaling-related pathways are also prominent in the liver and kidney, consistent with their dependence on receptor-mediated regulation and intercellular communication. Immune-related pathways appear across multiple tissues, particularly in those at environmental interfaces such as the skin and lung, underscoring general defense mechanisms. Neuronal pathways are uniquely enriched in the brain, as expected, with minor representation in reproductive tissues such as the ovary and prostate, possibly reflecting shared neuroendocrine signaling. Categories related to apoptosis, cell cycle, and transport show lower but consistent activation across all tissues, corresponding to baseline cellular maintenance and turnover processes. Overall, these distributions indicate that the model captures biologically coherent, tissue-specific pathway preferences, supporting the interpretability of its learned representations.

Importantly, the pathways identified by the ablation-based analysis are not generic; rather, they correspond to tissue-relevant biological processes. In the brain, top-ranked pathways included immune signaling and apoptosis-related mechanisms such as FOXO-mediated transcription of cell-death genes and complement-associated cascades. These processes are known to contribute to neuroinflammation and neural maintenance, suggesting that immune–neuronal interactions play a meaningful role in tissue-specific predictive behavior.

In breast tissue, pathways related to FGF signaling, mTOR signaling, and estrogen-regulated transcription were prioritized—consistent with hormone-driven growth and metabolic regulation central to mammary physiology and cancer susceptibility. Mitochondrial and glycosylation pathways were also prominent, reflecting high metabolic turnover and protein modification demands. For the liver, a key organ in xenobiotic metabolism, the model emphasized FGFR1/2 signaling, nitric oxide pathways, and xenobiotic detoxification, alongside lipid and bile acid metabolism. These patterns align well with established hepatic functions in metabolic regulation and detoxification. Similarly, in the kidney, enriched pathways included heme scavenging, xenobiotic metabolism, and ion transport, reflecting its roles in toxin clearance and electrolyte balance. In cervical tissue, the analysis highlighted glycosylation, apoptosis, and lipid metabolism pathways—processes associated with epithelial differentiation and immune defense. Extending to the lung, ovary, and prostate, several shared pathways emerged, including ion channel regulation and mTOR–AKT signaling, suggesting conserved mechanisms of cellular signaling and stress response across epithelial tissues.

Across all tissues, general processes such as protein methylation, cell-cycle regulation, and lipid remodeling were consistently among the top contributors, indicating that transcriptional control and metabolic balance broadly support predictive performance. Taken together, these results suggest that Expresso learns coherent, biologically grounded patterns across tissues, demonstrating that its predictions arise from interpretable functional mechanisms rather than surface-level associations.

## 6.2 Prioritizing Clinically Relevant Compounds

To assess how useful the model could be in real-world drug discovery, we tested whether it could prioritize known, approved drugs for the tissues they’re actually used in. For the liver, it highly ranked Sorafenib Keating (2017); Kane et al. (2006), a well-established treatment for liver cancer; In breast tissue, it identified Fulvestrant Ciruelos et al. (2014); Nathan & Schmid (2017), an estrogen blocker used for breast cancer; For the brain, it correctly elevated Donepezil Kertesz (2004), commonly prescribed for Alzheimer’s. Not only does the model recognize effective compounds—it also deprioritizes the ones that failed. It correctly downgraded Fialuridine McKenzie et al. (1995), a drug that showed early promise but was withdrawn due to liver toxicity, and Talampanel, which had limited benefit in neurological diseases.

These results show that the model captures subtle tissue-specific molecular signals that reflect real clinical outcomes. By integrating biologically structured context into its predictions, it supports more targeted, human-relevant therapeutic development. In practice, Expresso enables researchers to focus efforts on compounds with higher potential for human relevance, narrowing the gap between computational screening and biological validation. By grounding its predictions in both molecular structure and the molecular state of the target tissue, it supports more strategic experimental planning and helps avoid costly trial-and-error. This integration of context-aware modeling into real-world workflows represents a shift toward more precise, biology-driven approaches in drug development, enhancing both efficiency and clinical confidence.

## 7 Conclusions

In this work, we presented Expresso, a multi-modal framework that enhances molecular activity prediction by incorporating tissue-specific biological context. While most existing models treat tissues as flat input, Expresso is designed to reason over the molecular and cellular environment in which compounds act. By embedding both molecular structure and rich tissue-level transcriptional information into a unified predictive model, Expresso bridges a critical gap between in vitro data and the complex biological settings of human tissues.

We conduct an empirical evaluation across nine benchmarks involving thousands of molecules, demonstrating the value of leveraging tissue-specific biological information for activity prediction. Expresso consistently outperforms baseline architectures, with performance gains largely attributed to our novel tissue encoder. This encoder models gene–pathway interactions and captures sample-level variation within tissue types. By grounding the model in biological context, Expresso enables more accurate and generalizable predictions of compound activity across diverse human tissues.

## References

- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Roman Bresson, Giannis Nikolentzos, George Panagopoulos, Michail Chatzianastasis, Jun Pang, and Michalis Vazirgiannis. KAGNNs: Kolmogorov-arnold networks meet graph learning. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL <https://openreview.net/forum?id=03UB1MCAMr>.
- Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491*, 2021.
- Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885*, 2020.
- Eva Ciruelos, Tomás Pascual, María Luisa Arroyo Vozmediano, Marta Blanco, Luis Manso, Lucía Parrilla, Cesar Muñoz, Estela Vega, Monica Jackelin Calderón, Blanca Sancho, et al. The therapeutic role of fulvestrant in the management of patients with hormone receptor-positive breast cancer. *The Breast*, 23(3):201–208, 2014.
- GTEx Consortium. The gtex consortium atlas of genetic regulatory effects across human tissues. *Science*, 369(6509):1318–1330, 2020.
- David Croft, Gavin O’kelly, Guanming Wu, Robin Haw, Marc Gillespie, Lisa Matthews, Michael Caudy, Phani Garapati, Gopal Gopinath, Bijay Jassal, et al. Reactome: a database of reactions, pathways and biological processes. *Nucleic acids research*, 39(suppl\_1):D691–D697, 2010.
- Nathan C Frey, Ryan Soklaski, Simon Axelrod, Siddharth Samsi, Rafael Gomez-Bombarelli, Connor W Coley, and Vijay Gadeppally. Neural scaling of deep chemical models. *Nature Machine Intelligence*, 5(11):1297–1305, 2023.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- Yoshinobu Igarashi, Noriyuki Nakatsu, Tomoya Yamashita, Atsushi Ono, Yasuo Ohno, Tetsuro Urushidani, and Hiroshi Yamada. Open tg-gates: a large-scale toxicogenomics database. *Nucleic acids research*, 43(D1):D921–D927, 2015.
- John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative models for graph-based protein design. *Advances in neural information processing systems*, 32, 2019.
- Robert C Kane, Ann T Farrell, Haleh Saber, Shenghui Tang, Gene Williams, Josephine M Jee, Chengyi Liang, Brian Booth, Nallaperumal Chidambaram, David Morse, et al. Sorafenib for the treatment of advanced renal cell carcinoma. *Clinical cancer research*, 12(24):7271–7278, 2006.
- Minoru Kanehisa and Susumu Goto. Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids research*, 28(1):27–30, 2000.
- Gillian M Keating. Sorafenib: a review in hepatocellular carcinoma. *Targeted oncology*, 12(2):243–253, 2017.
- Andrew Kertesz. Long-term donepezil treatment in patients with alzheimer’s disease. *Current Neurology and Neuroscience Reports*, 4(6):459–460, 2004.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.
- Juncai Li and Xiaofei Jiang. Mol-bert: An effective molecular representation with bert for molecular property prediction. *Wireless Communications and Mobile Computing*, 2021(1):7181815, 2021.

- Qiao Liu, Zhiqiang Hu, Rui Jiang, and Mu Zhou. Deepcdr: a hybrid graph convolutional network for predicting cancer drug response. *Bioinformatics*, 36(Supplement\_2):i911–i918, 2020.
- Xiaofan Liu, Yuhuan Tao, Zilin Cai, Pengfei Bao, Hongli Ma, Kexing Li, Mengtao Li, Yunping Zhu, and Zhi John Lu. Pathformer: a biological pathway informed transformer for disease diagnosis and prognosis using multi-omics data. *Bioinformatics*, 40(5):btac316, 2024.
- Polina Mamoshina, Armando Vieira, Evgeny Putin, and Alex Zhavoronkov. Applications of deep learning in biomedicine. *Molecular pharmaceuticals*, 13(5):1445–1454, 2016.
- Robin McKenzie, Michael W Fried, Richard Sallie, Hari Conjeevaram, Adrian M Di Bisceglie, Yoon Park, Barbara Savarese, David Kleiner, Maria Tsokos, Carlos Luciano, et al. Hepatic failure and lactic acidosis due to fialuridine (fiau), an investigational nucleoside analogue for chronic hepatitis b. *New England Journal of Medicine*, 333(17):1099–1105, 1995.
- Mark R Nathan and Peter Schmid. A review of fulvestrant in breast cancer. *Oncology and therapy*, 5(1):17–29, 2017.
- Darryl Nishimura. Biocarta. *Biotech Software & Internet Report: The Computer Software Journal for Scient*, 2(3):117–120, 2001.
- Christoph Ogris, Dimitri Guala, Thomas Helleday, and Erik LL Sonnhammer. A novel method for crosstalk analysis of biological networks: improving accuracy of pathway annotation. *Nucleic acids research*, 45(2):e8–e8, 2017.
- Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Grover: Self-supervised message passing transformer on large-scale molecular data. *arXiv preprint arXiv:2007.02835*, 2(3):17, 2020.
- Carl F Schaefer, Kira Anthony, Shiva Krupa, Jeffrey Buchoff, Matthew Day, Timo Hannay, and Kenneth H Buetow. Pid: the pathway interaction database. *Nucleic acids research*, 37(suppl\_1):D674–D679, 2009.
- Thomas Schmitt, Christoph Ogris, and Erik LL Sonnhammer. Funcoup 3.0: database of genome-wide functional coupling networks. *Nucleic acids research*, 42(D1):D380–D388, 2014.
- Philipp Seidl, Andreu Vall, Sepp Hochreiter, and Günter Klambauer. Enhancing activity prediction models in drug discovery with the ability to understand human language, 2023. URL <https://arxiv.org/abs/2303.03363>.
- Parthasarathy Suryanarayanan, Yunguang Qiu, Shreyans Sethi, Diwakar Mahajan, Hongyang Li, Yuxin Yang, Elif Eyigoz, Aldo Guzman Saenz, Daniel E Platt, Timothy H Rumbell, et al. Multi-view biomedical foundation models for molecule-target and property prediction. *arXiv preprint arXiv:2410.19704*, 2024.
- Daniel L Svoboda, Trey Saddler, and Scott S Auerbach. An overview of national toxicology program’s toxicogenomic applications: Drugmatrix and toxfx. *Advances in computational toxicology: Methodologies and applications in regulatory science*, pp. 141–157, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Fang Wu, Dragomir Radev, and Stan Z Li. Molformer: Motif-based transformer on 3d heterogeneous molecular graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 5312–5320, 2023.
- You Wu, Qiao Liu, Yue Qiu, and Lei Xie. Deep learning prediction of chemical-induced dose-dependent and context-specific multiplex phenotype responses and its application to personalized alzheimer’s disease drug repurposing. *PLoS computational biology*, 18(8):e1010367, 2022.