

# GALAX: GRAPH-AUGMENTED LANGUAGE MODEL FOR EXPLAINABLE REINFORCEMENT-GUIDED SUBGRAPH REASONING IN PRECISION MEDICINE

Heming Zhang<sup>1</sup>, Di Huang<sup>2</sup>, Wenyu Li<sup>2</sup>, Michael Province<sup>3</sup>

Yixin Chen<sup>2</sup>, Philip Payne<sup>1</sup>, Fuhai Li<sup>1,2\*</sup>

<sup>1</sup>I2DB, <sup>2</sup>Dept. of CSE, <sup>3</sup>Dept. of Genetics, Washington University in St. Louis

\*Correspondence: fuhai.li@wustl.edu

## ABSTRACT

In precision medicine, quantitative multi-omic features, topological context, and textual biological knowledge play vital roles in identifying disease-critical signaling pathways and targets, guiding the discovery of novel therapeutics and effective treatment strategies. Existing pipelines capture only one or two of these, thereby limiting mechanistic interpretability. Although Process Reward Models (PRMs) aim to guide reasoning in LLMs, they remain limited by coarse step definitions, unreliable intermediate evaluation, and vulnerability to reward hacking with added computational cost. These gaps motivate jointly integrating quantitative multi-omic signals, topological structure with node annotations, and literature-scale text via LLMs, using subgraph reasoning as the principle bridge linking numeric evidence, topological knowledge and language context. To resolve this challenge, we propose **GALAX** (Graph Augmented LAnguage model with eXplainability), an innovative framework that integrates pretrained Graph Neural Networks (GNNs) into Large Language Models (LLMs) via reinforcement learning guided by a Graph Process Reward Model (GPRM), which generates disease-relevant subgraphs in a step-wise manner initiated by an LLM and iteratively evaluated by a pretrained GNN and schema-based rule check, enabling process-level supervision without explicit labels. As an application, we also introduced **Target-QA**, a benchmark combining CRISPR-identified targets, multi-omic profiles, and biomedical graph knowledge across diverse cancer cell lines, which enables GNN pretraining for supervising step-wise graph construction and supports long-context reasoning over text-numeric graphs (TNGs), providing a scalable and biologically grounded framework for explainable, reinforcement-guided subgraph reasoning toward reliable and interpretable target and pathway discovery in precision medicine.

## 1 INTRODUCTION

Identifying therapeutic targets and elucidating disease mechanisms are main challenges in precision medicine (Steyaert et al., 2023; Topol, 2019). CRISPR-based gene editing has revolutionized functional genomic by enabling high-throughput perturbation of gene function across diverse cellular contexts (Shalem et al., 2014; Li et al., 2023). In oncology, large-scale CRISPR screens in cancer cell lines and patient-derived models have revealed context-specific genetic vulnerabilities, providing a robust experimental foundation for biomarker and target discovery (Shi et al., 2015). Despite these advances, computationally predicting key targets from multi-omic profiles and interpreting their mechanistic role in disease progression remains difficult. In particular, bridging omic data with interpretable explanations of molecular mechanisms continues to be a critical unmet need (Zhang et al., 2024a). Traditional approaches, such as differential expression analysis or essentiality scoring, lack the capacity to model the hierarchical and cross-modal dependencies in molecular networks, often overlooking key regulatory redundancies and pathway-level dynamics. Recent graph-based models have shown promise in outcome prediction tasks (Ren et al., 2024), yet they typically lack the structured supervision necessary for accurate target prioritization and mechanism discovery and seldom jointly integrate quantitative multi-omic features, topological structure with node annotations, and literature-scale text—limiting mechanistic interpretability.

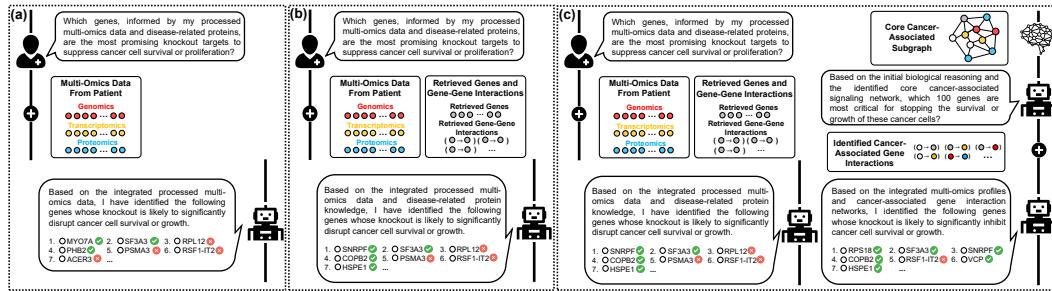
Meanwhile, Large Language Models (LLMs) have demonstrated strong capabilities in natural language understanding and reasoning, particularly through techniques like in-context learning (ICL) (Brown et al., 2020) and chain-of-thought (CoT) prompting (Wei et al., 2022), which enable multi-step reasoning. However, LLMs often suffer from hallucination and lack grounding in structured knowledge, especially in scientific domains. To mitigate these issues, Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) and its graph-based variants, such as RoG (Luo et al., 2023), SubgraphRAG (Li et al., 2024), GNN-RAG (Mavromatis & Karypis, 2025) and G-Retriever (He et al., 2024), have been proposed to enhance LLM performance by incorporating external knowledge graph. Despite their utility, these approaches still focus on final answer accuracy and give little attention to the reliability of intermediate reasoning. The retrieved subgraphs are noisy, large, and lack ground-truth mechanistic structure, making supervised retrieval unstable. Most existing models also fail to integrate numerical omic signals, causing the loss of cell line-specific information needed for target discovery. On the other hand, the Process Reward Model (PRM) framework has been introduced to provide fine-grained supervision over intermediate steps in reasoning tasks (Luo et al., 2024; Lightman et al., 2023; Uesato et al., 2022; Wang et al., 2023). PRMs provide step-wise supervision by assigning intermediate rewards to reinforcement learning (RL) agents, forming the foundation for Large Reasoning Models (LRMs) trained with Reinforcement Learning with Human Feedback (RLHF) (Bai et al., 2022), Proximal Policy Optimization (PPO) (Schulman et al., 2017), and Group Relative Policy Optimization (GRPO) (Shao et al., 2024). For example, StepGRPO (Zhang et al., 2025b) extends GRPO by incorporating rule-based step-wise rewards to supervise each intermediate reasoning step, addressing the sparse reward problem and enhancing multi-step reasoning in multimodal language models. However, PRMs face key limitations in defining fine-grained reasoning steps, verifying intermediate correctness, and hacking for model-based rewards (Gao et al., 2023), further complicating training.

These challenges are amplified in biomedicine: reasoning over multi-omic, gene-regulatory text-numeric graphs (TNGs) lacks ground-truth stepwise annotations, making intermediate supervision infeasible, and the combinatorial explosion of biological paths renders exhaustive planning or retrieval impractical. We propose **GALAX** (Graph-Augmented LAnguage model with eXplainability), which couples LLMs with a pretrained GNN under reinforcement learning guided by a **Graph Process Reward Model (GPRM)**. Instead of explicit labels, GALAX uses the GNN as a stepwise supervisor and schema-based rule term to check validity, scoring intermediate subgraphs (partial signaling cascades) for biological plausibility and cancer relevance to provide fine-grained, graph-based rewards. GALAX prompts an LLM to propose candidate targets from multi-omic profiles and partial knowledge graphs, then an RL graph generator assembles task-specific cancer subnetworks under GPRM scoring—translating language reasoning into interpretable graph construction and yielding mechanistically grounded, patient-specific subnetworks for target prioritization. To evaluate, we introduce **Target-QA**, a benchmark integrating multi-omic data, biomedical graph knowledge, and CRISPR screening outcomes across diverse cancer cell lines. Together, GALAX and Target-QA deliver a scalable, reinforcement-guided solution for interpretable, patient-specific target identification and disease-mechanism discovery.

## 2 RELATED WORK

**LLMs Augmented with Knowledge and Graph Structures** Prompt tuning has emerged as a lightweight and scalable method for adapting LLMs to downstream tasks without full finetuning (Korbak et al., 2023; Lester et al., 2021). While effective, it operates over flat text representations and struggles to incorporate structured domain knowledge or multi-modal signals. To address this, Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) and graph-augmented approaches such as G-Retriever (He et al., 2024), RoG (Luo et al., 2023), SubgraphRAG (Li et al., 2024) and GNN-RAG (Mavromatis & Karypis, 2025) have been proposed. However, these methods depend on accurate subgraph retrieval and still lack support for reliable reasoning. Most existing models also fail to integrate numerical omic signals, causing the loss of cell line-specific information needed for target discovery and leaving them poorly suited for large, patient-specific text-numeric graphs such as multi-omic signaling networks.

**Reinforcement Learning for Step-wise Reasoning** Reinforcement learning has been instrumental in aligning LLM behavior through methods like RLHF (Bai et al., 2022), PPO (Schulman et al., 2017), and GRPO (Shao et al., 2024). The Process Reward Model (PRM) (Luo et al.,



**Figure 1: Representative paradigms of patient-specific target prediction in-context LLM workflows and comparison with our approach.** (a) User-specific in-context prompts with basic query input. (b) Augmentation with retrieved knowledge relevant to the query. (c) GALAX framework: combining user-specific prompts, biomedical knowledge, and reinforcement-guided subgraph reasoning for explainable target prediction.

2024) enables step-wise supervision and has been adopted in Large Reasoning Models (LRMs). However, PRMs face key challenges: fine-grained step definitions are ambiguous, intermediate correctness is hard to validate, and model-based rewards can lead to reward hacking (Gao et al., 2023). These challenges are particularly acute in biomedical domains, where reasoning is inherently unstructured and lacks step-wise annotations, making conventional PRM pipelines impractical.

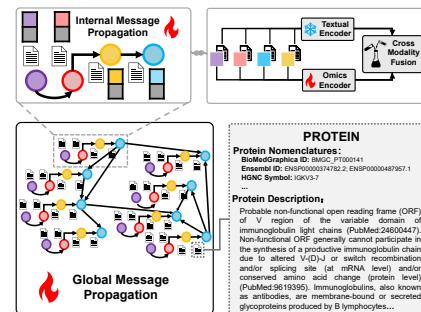
**Multi-omic Data Integration in Biomedical AI** From a biological standpoint, the integration of genomic, transcriptomic, and proteomic has been essential for understanding disease mechanisms and therapeutic vulnerabilities (Hasin et al., 2017; Kristensen et al., 2014). Traditional approaches rely on statistical fusion or dimensionality reduction (Meng et al., 2016; Shen et al., 2009; Rohart et al., 2017; Argelaguet et al., 2018; Nguyen & Wang, 2020), which overlook the hierarchical and interconnected nature of molecular data. More recently, GNN-based models like MOGONET (Wang et al., 2021) and MoGCN (Li et al., 2022) have demonstrated the value of structured graph reasoning for cancer subtype classification and biomarker identification. However, these models are primarily designed for outcome prediction and often fall short in identifying actionable biomarkers associated with specific disease mechanisms.

GALAX addresses the limitations of existing models by introducing a RL-guided framework that dynamically constructs biologically relevant subgraphs for each patient or cell line. This enables interpretable, context-sensitive target prioritization that adapts to both multi-omic features and disease-specific graph as the text-numeric format. To our knowledge, GALAX is the first to unify numerical multi-omic signals, literature-scale textual information, and biological topology under a reinforcement learning paradigm with biologically grounded supervision, which learns to reason through step-wise subgraph generation, guided by an authoritative biomedical GRPM.

### 3 PROBLEM FORMULATION

Identifying key targets and uncovering disease mechanisms remains a major challenge in precision medicine. To address this, we adopt the Text-Omic Signaling Graph (TOSG) (Zhang et al., 2025a), which integrates multi-omic features and biomedical context into a unified graph structure. Built upon TOSG, our model **GALAX** couples LLM-based hypothesis generation with reinforcement-guided subgraph reasoning to enable interpretable, patient-specific target predictions on the Target-QA benchmark.

**TOSG Construction** We construct the TOSG by integrating three modalities: numerical omic evidence  $\mathcal{X}^{(0)}$ , textual entity descriptions  $\mathcal{T}$ , and topological information within a unified biomedical knowledge graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ . Here,  $\mathcal{V}$  represents the topological entities and  $\mathcal{E}$  denotes the biological relations. The detailed formulation of each modality is presented as follows. **Feature Space:** Specific of numerical evidence (Dempster et al., 2019), we formul



**Figure 2:** Demonstration of TOSG Construction and Graph Foundation Model

ally, utilizing the DepMap as the source, we can generate this raw data as  $\mathcal{X}^{(0)} = \{X_n^{(0)}\}_{n=1}^{N^{(0)}}$ ,

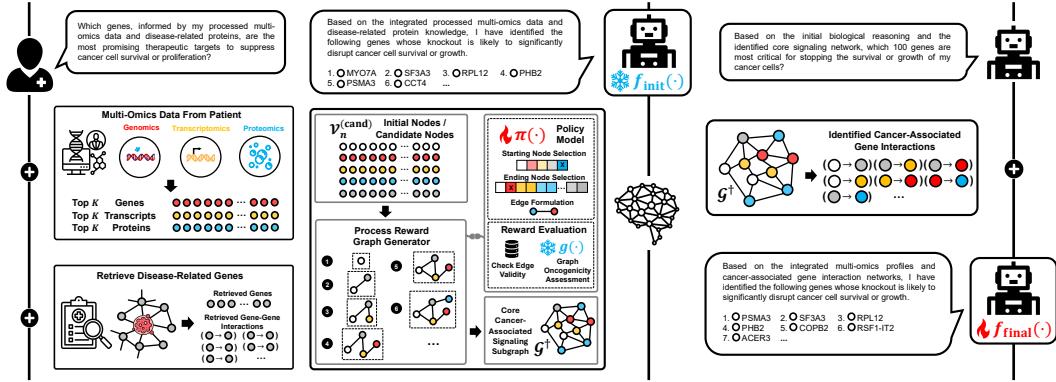
where  $X_n^{(0)} \in \mathbb{R}^M$  comprises of  $N^{(0)}$  samples with  $M$  entities. These samples are paired with binary labels for cancerous or non-cancerous cell lines, denoted as  $\mathcal{Y}^{(0)} \in \{0, 1\}^{N^{(0)}}$ . **Textual Space:** Each entity in the TOSG is associated with a name and description, represented as  $\mathcal{T} = \{T_{\text{name}}, T_{\text{desc}}\}$ , where  $|T_{\text{name}}| = |T_{\text{desc}}| = M$ . **Topological Space:** The  $M$  entities in each sample are composed of multiple components derived from genomic, transcriptomic, and proteomic contexts. These multi-omic features are integrated into the TOSG using an existing integration tool, BioMedGraphica (Zhang et al., 2024b), resulting in a text-attributed knowledge graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ . The set of vertices is defined as  $\mathcal{V} = \{\mathcal{V}^{(pm)}, \mathcal{V}^{(g)}, \mathcal{V}^{(t)}, \mathcal{V}^{(p)}\}$ , representing promoter, gene, transcript, and protein entities, respectively. The size of each set is given by  $|\mathcal{V}^{(pm)}| = m^{(pm)}$ ,  $|\mathcal{V}^{(g)}| = m^{(g)}$ ,  $|\mathcal{V}^{(t)}| = m^{(t)}$ , and  $|\mathcal{V}^{(p)}| = m^{(p)}$ , such that  $|\mathcal{V}| = m^{(pm)} + m^{(g)} + m^{(t)} + m^{(p)} = M$ . Correspondingly, we map the cell-specific omic features for the  $n$ -th sample, denoted as  $X_n^{(0)} \in \mathbb{R}^M$ , using a compact vector concatenation of the four entity types:  $X_n^{(0)} = \left[ \mathbf{x}_n^{(pm)} \oplus \mathbf{x}_n^{(g)} \oplus \mathbf{x}_n^{(t)} \oplus \mathbf{x}_n^{(p)} \right]$ , where  $\mathbf{x}_n^{(\cdot)}$  represents the feature vector for the respective modality (e.g.,  $\mathbf{x}_n^{(p)} \in \mathbb{R}^{m^{(p)}}$  corresponds to protein levels). In detail, this graph can be decomposed into two subgraphs:  $\mathcal{G}^{(\text{in})} = (\mathcal{V}^{(\text{in})}, \mathcal{E}^{(\text{in})})$  and  $\mathcal{G}^{(\text{PPI})} = (\mathcal{V}^{(\text{PPI})}, \mathcal{E}^{(\text{PPI})})$ . Here,  $\mathcal{G}^{(\text{in})}$  captures the internal signaling processes for protein translation. As shown in Figure 2, internal propagation follows the central dogma (Crick, 1970): promoter (**purple**)  $\rightarrow$  gene (**red**)  $\rightarrow$  transcript (**yellow**)  $\rightarrow$  protein (**blue**), with  $\mathcal{V} = \mathcal{V}^{(\text{in})}$  and  $|\mathcal{V}^{(\text{in})}| = M$ . Meanwhile,  $\mathcal{G}^{(\text{PPI})}$  represents the gene regulatory network structured around protein-protein interactions (PPI), where  $\mathcal{V}^{(\text{PPI})} = \mathcal{V}^{(p)}$ . In summary, the TOSG unifies numerical omic evidence, textual descriptions, and topological information into a Text-Numeric Graph, defined as  $\mathcal{G} = \{\mathcal{X}^{(0)}, \mathcal{T}, \mathcal{V}, \mathcal{E}\}$ .

**Target-QA Generation** We construct the feature set  $\mathcal{X}$  by filtering for cancer cell lines that contain both comprehensive annotations and curated CRISPR-based target information, defining the collective multi-omic evidence as set  $\mathcal{X} = \{X_n\}_{n=1}^N$ , where  $X_n \in \mathbb{R}^M$ . Subsequently, to align this quantitative data with the reasoning capabilities of LLMs, we systematically structure the input context for each instance by integrating textual metadata, molecular profiles, and interaction graphs, detailed as follows. **Omics Information:** To augment context with patient-specific data, we incorporate the top- $K$  features for the LLM. Due to input token constraints of LLMs (Achiam et al., 2023) and the presence of CpG sites with methylation beta values saturated at 1, which dominate rankings and render top- $K$  selection uninformative, we derive a concise multi-omic representation  $X_n^{(K)} = [\mathbf{g}_n^{(K)} \oplus \mathbf{t}_n^{(K)} \oplus \mathbf{p}_n^{(K)}]$  by extracting the top  $K$  features from genomic, transcriptomic, and proteomic modalities. Here,  $\mathbf{g}_n^{(K)}$ ,  $\mathbf{t}_n^{(K)}$ , and  $\mathbf{p}_n^{(K)}$  denote the ordered lists of selected gene, transcript, and protein names, respectively, which correspond directly to the omic-related graph nodes  $\mathcal{V}_n^{(\text{omic})}$ . **Sample Information:** Disease entity has an associated name and textual description with  $\mathcal{S} = \{S_{\text{name}}, S_{\text{desc}}\}$  where  $|\mathcal{S}| = m^{(S)}$ . Leveraging annotations from DepMap about cell lines, we mapped the cell line and disease names in samples to form the sets  $\mathcal{C} = \{c_n\}_{n=1}^N$  and  $\mathcal{S}' = \{s'_n\}_{n=1}^N$ , respectively. **Disease-related Protein Subgraph:** For cell line,  $c_n$ , we use the cell line related disease entity  $s'_n$  to retrieve from BioMedGraphica’s disease-target interaction graph  $\mathcal{G}^{(\text{DTI})} = \{\mathcal{V}^{(\text{DTI})}, \mathcal{E}^{(\text{DTI})}\}$ , where  $\mathcal{V}^{(\text{DTI})} = \{\mathcal{V}^{(S)}, \mathcal{V}^{(p)}\}$  includes disease and protein nodes. To provide structured graph context, we apply a subgraph retrieval strategy that extracts disease-relevant protein entities  $\mathcal{V}_n^{(p)} \subset \mathcal{V}^{(p)}$  and their interactions  $\mathcal{E}_n^{(\text{DTI})}$ , along with  $h$ -hop protein neighbors  $\mathcal{V}_n^{(h)}$  by extracting interactions from  $\mathcal{E}^{(\text{PPI})}$ , where  $\mathcal{V}_n^{(h)} \subset \mathcal{V}^{(p)}$ . This yields a sample-specific subgraph  $\mathcal{G}_n^{(\text{sub})} = \{\mathcal{V}_n^{(\text{sub})}, \mathcal{E}_n^{(\text{sub})}\}$ , where  $\mathcal{V}_n^{(\text{sub})} = \mathcal{V}_n^{(p)} \cup \mathcal{V}_n^{(h)}$ , with the full graph set  $\mathcal{G}^{(\text{sub})} = \{\mathcal{G}_n^{(\text{sub})}\}_{n=1}^N$ . Hence, each query  $Q_n = \{c_n, s'_n, X_n^{(K)}, \mathcal{G}_n^{(\text{sub})}\}$  is paired with answer  $A_n$  describing top- $\gamma$  CRISPR targets  $R_n = \{r_{n,1}, \dots, r_{n,\gamma}\}$ , yielding instance  $D_n = (Q_n, A_n)$  in dataset  $\mathcal{D} = \{D_n\}_{n=1}^N$ , stratified by TCGA (Weinstein et al., 2013) types (e.g.,  $\mathcal{D}_{\text{LUAD}}, \mathcal{D}_{\text{BRCA}}$ ).

**Patient-Specific Target Prediction with Explainability** GALAX is designed to generate not only accurate but also interpretable predictions by explicitly modeling the reasoning process through subgraph construction (see Figure 3). Given a query  $Q_n$ , the model produces both a prioritized target list  $\hat{A}_n$  and an explanatory subgraph  $\mathcal{G}_n^\dagger$ :

$$\hat{A}_n, \mathcal{G}_n^\dagger = f(Q_n, X_n, \mathcal{T}, \mathcal{E}; \theta_G, \theta_L) \quad (1)$$

The function  $f$  is composed of three modules (see Section 4.2): (1) an initial language model  $f_{\text{init}}$  that performs coarse reasoning and extracts candidate entities via initial answering; (2) a reinforcement-



**Figure 3: An overview of the GALAX workflow.** Given processed multi-omic profiles (genomic, transcriptomic, proteomic), a subgraph is retrieved by identifying disease-associated proteins and their  $h$ -hop neighbors. Then  $f_{\text{init}}(\cdot)$  proposes initial targets, refined by a reinforcement-guided graph generator  $\pi(\cdot)$  supervised by a pretrained graph foundation model  $g(\cdot)$ . The final subgraph,  $\mathcal{G}^\dagger$ , combined with the query, is passed to a second-stage LLM  $f_{\text{final}}(\cdot)$  for target prediction. The full pipeline enables explainable, patient-specific reasoning grounded in molecular biology and CRISPR evidence.

based graph generator  $\pi(\cdot)$  that incrementally constructs the explainable subgraph  $\mathcal{G}_n^\dagger$  under the guidance of a pretrained graph classifier  $g(\cdot)$  parameterized by  $\theta_G$ , using step-wise biological plausibility rewards; and (3) a final language model  $f_{\text{final}}(\cdot)$  that refines the prediction by reasoning over both the initial output and the generated subgraph context. The subgraph  $\mathcal{G}_n^\dagger$  serves as a transparent rationale, offering deeper insights into the disease mechanism.

## 4 GALAX

### 4.1 FOUNDATION MODELS PRETRAINING

**LLM Pretraining** We pretrain a large language model, denoted as  $f_L^{\text{pre}}$  with parameters  $\theta_L^{\text{pre}}$ , using curated text corpora. The input data comprises omic entity descriptions  $\mathcal{T}$ , disease annotations  $\mathcal{S}$ , protein–protein interactions  $\mathcal{E}^{(\text{PPI})}$ , and disease–target relationships  $\mathcal{E}^{(\text{DTI})}$ . As illustrated in Figure 5, this pretraining phase equips the model with foundational knowledge of biomedical terminology and relational structure (see Appendix C.1), thereby enhancing its capacity for downstream reasoning in biomedical tasks. Then we will continue pretraining the language model  $f_{\text{init}}$  parameterized by  $\theta_{\text{init}}$ , which are detailed in Appendix C.2.

**Graph Foundation Model Pretraining** We pretrain the graph encoder  $\theta_G^{\text{pre}}$  via a two-stage pipeline (see Figure 2). In stage one, a unified graph–language model  $f_G^{\text{pre}}$  is trained over node attributes  $\mathcal{X}^{(0)}$ , textual features  $\mathcal{T}$ , and edge set  $\mathcal{E}$ , with protein–protein interaction edges stochastically masked as  $\mathcal{E}_{\text{mask}} \sim \text{Bernoulli}(p)$ ,  $p < 1$ . The resulting representation is

$$\mathcal{H}^{\text{pre}} = f_G^{\text{pre}}(\mathcal{X}^{(0)}, \mathcal{T}, \mathcal{E}, \mathcal{E}_{\text{mask}}) \quad (2)$$

, where  $\mathcal{H}^{\text{pre}} \in \mathbb{R}^{N^{(0)} \times M \times d^{\text{pre}}}$  encodes contextualized entity states. In details, we generate edge mask  $\mathcal{E}_{\text{mask}} \sim \text{Bernoulli}(p)$ , where  $p < 1$  is the ratio of the masked edges for  $\mathcal{E}^{(\text{PPI})}$  to mask out the signaling flows in protein–protein interactions. Then, we apply internal message propagation with

$$\mathcal{H}_{\text{in}}^{\text{pre}} = \text{GNN}_{\text{in}}^{\text{pre}}(\text{ENC}_{\text{cross}}^{\text{pre}}(\mathcal{X}^{(0)}, \mathcal{T}), \mathcal{E}^{(\text{in})}) \quad (3)$$

, where  $\text{ENC}_{\text{cross}}^{\text{pre}}$  is a cross-modal encoder to align textual and omic features and  $\mathcal{H}_{\text{in}}^{\text{pre}} \in \mathbb{R}^{N^{(0)} \times M \times d^{(\text{in})}}$ . The first-stage pretraining captures gene regulatory patterns by performing masked global message passing over  $\mathcal{E}^{(\text{PPI})}$  by

$$\mathcal{H}^{\text{pre}} = \text{GNN}_{\text{PPI}}^{\text{pre}}(\mathcal{H}_{\text{in}}^{\text{pre}}, \mathcal{E}^{(\text{PPI})}, \mathcal{E}_{\text{mask}}) \quad (4)$$

In the second stage, downstream model  $f_G$  is initialized with pretrained parameters  $\theta_G^{\text{pre}}$  and is used to predict disease types from multi-omic inputs. The predicted class for each sample is given by:

$$\hat{\mathcal{Y}}^{(0)} = \arg \max_{o \in \mathcal{O}} \text{Softmax} \left[ \text{MLP}_G \left( f_G(\mathcal{X}^{(0)}, \mathcal{T}, \mathcal{E}; \theta_G^{\text{pre}}) \right) \right] \quad (5)$$

, where  $\mathcal{O}$  denotes the set of disease types, and  $f_G$  consists of the same architecture as  $\text{ENC}_{\text{cross}}$ ,  $\text{GNN}_{\text{in}}$ , and  $\text{GNN}_{\text{PPI}}$ . The model is trained to minimize the cross-entropy loss between the predicted probability distribution by  $\text{Softmax}[\text{MLP}_G(f_G(\cdot))]$  and the ground-truth label. The pretrained graph foundation model  $f_G$  serves as a structural proxy to guide step-wise reasoning in downstream tasks.

## 4.2 MODEL TRAINING

The pipeline of model couples initial answering, a reinforcement-guided subgraph generator, and final answering. Given the structured prompt  $P_n^{(\text{init})}$  from query  $Q_n$ , the pretrained model  $f_{\text{init}}$  outputs  $A_n^{(\text{init})}$  and a biomedical named entity recognition (NER)  $\phi$  extracts entities  $R_n^{(\text{init})}$  and maps them to proteins  $\mathcal{V}_n^{(\text{init})}$ , which will contribute to forming the start set  $\mathcal{V}_n^{(\text{start})}$  depending on situation. Node features are embedded by pretrained graph foundation model to obtain  $\mathcal{H}_{\text{in}}$  and formed  $X_n^{(\text{cand})}$ . Subgraph construction is framed as reinforcement learning problem consisting of four elements: **state**  $\mathcal{G}_n^{(i)} = (\mathcal{V}_n^{(i)}, \mathcal{E}_n^{(i)})$ ; **action**  $\Delta_n^{(i)} = (v_{\text{src}}^i, v_{\text{tgt}}^i)$  adds a single edge under feasibility masks; **policy** uses a message propagation (MSG) module to produce  $X_n^{(i)} = \pi_{\text{MSG}}(\mathcal{G}_n^{(i)}, X_n^{(\text{cand})})$  and two masked probability function  $\pi_{\text{SRC}}, \pi_{\text{TGT}}$  to sample  $v_{\text{src}}^i$  and  $v_{\text{tgt}}^i$ ; **reward** combines feedback from a pretrained classifier  $g(\cdot)$ , a rollout averaging  $L$  simulated continuations, and a rule-based term  $\mathcal{R}_{\text{rule}}$  that penalizes schema violations. We accept an action  $e$  only when  $\mathcal{R}_{\text{total}}^{(i)} > 0$  and update the generator with reward-weighted cross-entropy with early stopping; the best subgraph  $\mathcal{G}_n^{\dagger}$  is retained in  $\Omega$  stochastic run. For final answering,  $\mathcal{G}_n^{\dagger}$  is verbalized in expert mode (Fatemi et al., 2023) and appended to  $Q_n$  to form  $P_n^{(\text{final})}$ ; the model  $f_{\text{final}}$  is finetuned with token-level cross-entropy against  $A_n$ . What follows is a detailed exposition of the model design and training framework.

**Initial Answering** With the pretrained language model  $f_{\text{init}}$  based on  $f_L$ , parameterized by  $\theta_{\text{init}}$ , the input to the model is a structured prompt  $P_n^{(\text{init})}$  derived from the original query  $Q_n$ , specifically for turning  $\mathcal{G}_n^{(\text{sub})}$  into a graph expert format (see Appendix C.2 for details), and it will output  $A_n^{(\text{init})}$ . Afterwards, an NER function,  $\phi$ , is applied to extract biomedical entities from  $A_n^{(\text{init})}$  by  $R_n^{(\text{init})} = \phi(A_n^{(\text{init})})$ , where  $R_n^{(\text{init})} = \{r_{n,1}^{(\text{init})}, r_{n,2}^{(\text{init})}, \dots, r_{n,\alpha}^{(\text{init})}\}$ . And which are then mapped to corresponding protein nodes,  $\mathcal{V}_n^{(\text{init})} = \{v_{n,\text{init},1}^{(p)}, v_{n,\text{init},2}^{(p)}, \dots, v_{n,\text{init},\alpha'}^{(p)}\}$ .

**Process Reward Graph Generator** The initial node set  $\mathcal{V}_n^{(\text{start})}$  is selected based on a predefined priority: if the disease-related protein set  $\mathcal{V}_n^{(p)}$  is available, the top  $\eta$  most relevant entities are used; if not, the top  $\eta$  entities from the initialization set  $\mathcal{V}_n^{(\text{init})}$  are selected. If both are unavailable,  $\eta$  nodes are randomly sampled from the omic-derived set  $\mathcal{V}_n^{(\text{omic})}$ . Formally,

$$\mathcal{V}_n^{(\text{start})} = \begin{cases} \text{Top-}\eta(\mathcal{V}_n^{(p)}), & \text{if } \mathcal{V}_n^{(p)} \neq \emptyset \\ \text{Top-}\eta(\mathcal{V}_n^{(\text{init})}), & \text{else if } \mathcal{V}_n^{(\text{init})} \neq \emptyset \\ \text{Sample-}\eta(\mathcal{V}_n^{(\text{omic})}), & \text{otherwise} \end{cases} \quad (6)$$

The candidate set was generated based on  $\mathcal{V}_n^{(\text{cand})} = \mathcal{V}_n^{(\text{init})} \cup \mathcal{V}_n^{(\text{sub})} \cup \mathcal{V}_n^{(\text{omic})}$  and  $\mathcal{V}_n^{(\text{start})} \subset \mathcal{V}_n^{(\text{cand})}$ . And the features of  $\mathcal{X}$  are precomputed using a graph encoder  $f_G^{\text{pre}}$  with parameters  $\theta_G^{\text{pre}}$ , followed by pretrained modules  $\text{ENC}_{\text{cross}}$  and  $\text{GNN}_{\text{in}}$  with parameters  $\theta_{\text{cross}}^G$  and  $\theta_{\text{in}}^G$ , respectively:

$$\mathcal{H}_{\text{in}} = \text{GNN}_{\text{in}} \left( \text{ENC}_{\text{cross}} \left( f_G^{\text{pre}}(\mathcal{X}, \mathcal{T}, \mathcal{E}; \theta_G^{\text{pre}}), \mathcal{T}; \theta_{\text{cross}}^G \right), \mathcal{E}^{(\text{in})}; \theta_{\text{in}}^G \right) \in \mathbb{R}^{N \times M \times d^{(\text{in})}} \quad (7)$$

The candidate node features  $X_n^{(\text{cand})}$  are selected from the precomputed representation  $\mathcal{H}_{\text{in}}$ . Then we introduce a reinforcement-guided graph generator, denoted as  $\pi(\cdot)$ , where the policy operates over sample-specific graph states and candidate sets, enabling personalized subgraph construction. At step  $i$ , the current graph state is defined as  $\mathcal{G}_n^{(i)} = \{\mathcal{V}_n^{(i)}, \mathcal{E}_n^{(i)}\}$  after applying the action of constructing edge  $(v_{\text{src}}^i, v_{\text{tgt}}^i)$ . The next graph state  $\mathcal{G}_n^{(i+1)}$  is formed by adding an edge between a sampled source node  $v_{\text{src}}^i$  and target node  $v_{\text{tgt}}^i$ . To compute the probabilities of selecting these nodes, we embed the node features from  $\mathcal{G}_n^{(i)}$  together with the candidate features  $X_n^{(\text{cand})}$  as:

$$X_n^{(i)} = \pi_{\text{MSG}}(\mathcal{G}_n^{(i)}, X_n^{(\text{cand})}) \quad (8)$$

, where  $\pi_{\text{MSG}}$  consists of a message propagation (MSG) module and a feature selection mechanism that extracts node embeddings corresponding to  $\mathcal{V}_n^{(i)}$  from the propagated features. Based on embeddings, the source and target nodes are sampled according to the generated probabilities with

$$v_{\text{src}}^i \sim \pi_{\text{SRC}}(X_n^{(i)}, \mathcal{M}_{\text{SRC}}); \quad v_{\text{tgt}}^i \sim \pi_{\text{TGT}}(X_n^{(i)}, \mathcal{M}_{\text{TGT}}; v_{\text{src}}^i) \quad (9)$$

, where  $\pi_{\text{SRC}}$  and  $\pi_{\text{TGT}}$  are returned with probability implemented via MLPs and softmax by masking  $\mathcal{M}_{\text{SRC}}$  and  $\mathcal{M}_{\text{TGT}}$ , which restrict source selection to nodes in  $\mathcal{G}_n^{(i)}$  and exclude the source when selecting the target. Based on the probability, nodes  $v_{\text{src}}^i$  and  $v_{\text{tgt}}^i$  will be selected by the sampling function. The selected node pair then forms the updated graph state  $\mathcal{G}_n^{(i+1)}$  for next step.

**Reinforcement-Guided Reward and Training** To guide the graph generation process, we define a rollout-based reward function that combines immediate classifier feedback with future trajectory simulation. At generation step  $i$ , the intermediate graph is represented as  $\mathcal{G}_n^{(i+1)} = \{\mathcal{V}_n^{(i+1)}, \mathcal{E}_n^{(i+1)}\}$  after applying the action of constructing edge  $(v_{\text{src}}^i, v_{\text{tgt}}^i)$ . We define  $g(\cdot)$  as a pretrained graph classifier that computes class probabilities by applying a GNN encoder  $\text{GNN}_{\text{PPI}}^G$  followed by a projection head  $\text{MLP}_G$ , formally expressed as  $g(\cdot) = \text{Softmax} [\text{MLP}_G (\text{GNN}_{\text{PPI}}^G(\cdot))]$ . The model is parameterized by pretrained weights  $\theta_{\text{PPI}}^G$  and  $\theta_{\text{MLP}}^G$ , and outputs a probability distribution over classes in  $\mathcal{O}$ . Let  $o^* \in \mathcal{O}$  denote the target class. The reward  $\mathcal{R}_n^{(i)}$  for step  $i$  is defined as:

$$\mathcal{R}_n^{(i)} = g_{o^*}(\mathcal{G}_n^{(i+1)}) - \frac{1}{|\mathcal{O}|} + \lambda \cdot \frac{1}{L} \sum_{\ell=1}^L \left[ g_{o^*}(\text{Rollout}_\ell(\mathcal{G}_n^{(i+1)})) - \frac{1}{|\mathcal{O}|} \right] \quad (10)$$

Here,  $g_{o^*}(\mathcal{G})$  refers to the probability assigned to the target class  $o^*$ , and  $\text{Rollout}_\ell(\cdot)$  simulates the  $\ell$ -th full trajectory by continuing generation from the current partial graph using the current policy. The hyperparameter  $\lambda$  balances intermediate and future rollout-based feedback. To ensure reasoning aligning with biological plausibility, we incorporate a rule-based reward term  $\mathcal{R}_{\text{rule}}(\mathcal{G}_n^{(i+1)})$  that penalizes invalid edges according to relations from BioMedGraphica. The final reward is:

$$\mathcal{R}_{\text{total}}^{(i)} = \mathcal{R}_n^{(i)} + \lambda_{\text{rule}} \cdot \mathcal{R}_{\text{rule}}(\mathcal{G}_n^{(i+1)}) \quad (11)$$

This formulation guides the generation process toward subgraphs that are both predictive of the target class and consistent with domain-specific biological priors. And we used the greedy acceptance where if  $\mathcal{R}_{\text{total}}^{(i)} > 0$ , set  $\mathcal{G}_n^{(i+1)}$  as current state; otherwise keep the previous state. Then in each step  $i$ , the model will be trained with loss function,

$$\mathcal{L}_{\text{step}} = -\mathcal{R}_{\text{total}}^{(i)} [\text{CE}(v_{\text{src}}^i, \pi_{\text{SRC}}(X_n^{(i)}, \mathcal{M}_{\text{SRC}})) + \text{CE}(v_{\text{tgt}}^i, \pi_{\text{TGT}}(X_n^{(i)}, \mathcal{M}_{\text{TGT}}; v_{\text{src}}^i))] \quad (12)$$

, where the generator is optimized with reward-weighted cross-entropy (CE) function. In practice, we sample multiple candidate subgraphs under the policy parameterized by  $\theta_\pi$  across multiple runs  $\Omega$ , and select the optimal subgraph  $\mathcal{G}_n^\dagger$ .

**Final Answer Generation with Prompt Tuning** The optimal subgraph  $\mathcal{G}_n^\dagger$  is converted into a structured textual description via expert mode (see details in Appendix D.1), which will be appended to original query  $Q_n$  to form final graph-augmented prompt  $P_n^{(\text{final})}$ . With language model  $f_{\text{final}}$  based on pretrained  $f_{\text{init}}$ , it generates the output sequence  $\hat{A}_n = \{\hat{a}_{n,1}, \hat{a}_{n,2}, \dots, \hat{a}_{n,J'}\}$  according to:

$$\xi_{\theta_{\text{final}}}(\hat{A}_n \mid Q_n, \mathcal{G}_n^\dagger) = \prod_{j=1}^{J'} \xi_{\theta_{\text{final}}}(\hat{a}_{n,j} \mid \hat{a}_{n,< j}, P_n^{(\text{final})}) \quad (13)$$

To align the model output with the refined ground-truth answer  $A_n$ , which contains the top  $\gamma$  CRISPR-prioritized gene targets for sample  $n$ , we finetune the model by:

$$\mathcal{L}_{\text{final}} = - \sum_{n=1}^N \sum_{j=1}^{J'} \log \xi_{\theta_{\text{final}}}(\hat{a}_{n,j} \mid \hat{a}_{n,< j}, P_n^{(\text{final})}) \quad (14)$$

This objective encourages the model to internalize the structured reasoning encoded in  $\mathcal{G}_n^\dagger$  to generate biologically grounded answers. After generation, an NER function  $\phi$  will extract entities from the model's output  $\hat{A}_n$  with  $\hat{R}_n = \phi(\hat{A}_n) = \{\hat{r}_{n,1}, \hat{r}_{n,2}, \dots, \hat{r}_{n,\beta}\}$ . These predicted protein targets are used for evaluating biological relevance and overlap within reference targets in  $A_n$ .

## 5 EXPERIMENTS

**Datasets** We construct the Target-QA dataset over multiple TCGA cancer types using DepMap, integrating multi-omic features (epigenomic, genomic, transcriptomic, proteomic) and metadata from cancer cell lines. Each QA pair consists of an input query, including multi-omic and cell line information, and an output answer comprising the top- $\gamma$  ( $\gamma = 100$ ) CRISPR-prioritized targets. The final dataset contains 363 QA pairs from cancerous cell lines after integration and preprocessing. We use an 80/20 train-test split and repeat experiments across four randomized seeds to ensure stability and generalization. For foundation model pretraining, we collect dataset of 336 samples with multi-omic features from unannotated samples, including both disease and control groups (297 cancerous, 39 non-cancerous), and apply stratified sampling to address class imbalance. Full data processing and cohort composition are described in Appendix B.

**Experimental Setup** We initialize the language model with LLaMA3-8B-Instruct (Grattafiori et al., 2024), pretrained on biomedical terminology and curated textual descriptions involving protein–protein and disease–protein relationships from BioMedGraphica to enhance domain-specific vocabulary and biological context understanding (see Figure 5). For graph encoding, we use BioBERT-v1.1 (Lee et al., 2020) for text embeddings and Graph Attention Networks (GAT) (Veličković et al., 2017) to learn topological features from protein interaction graphs, incorporating random edge masking to improve robustness (see Figure 2 global message propagation). The pretrained GNN achieves 64.4% AUC in edge prediction, and 99.46% / 96.15% accuracy on disease type classification (train/test). Named entities are extracted using GPT-4o-mini via ChatGPT API (Hurst et al., 2024). GALAX is trained using the Adam optimizer on two NVIDIA H100 GPUs (80GB). We set the number of top omic features per modality to  $K=10$ , the maximum subgraph rollout depth to  $L=5$ , and the number of candidate starting nodes  $\eta=20$ . The reward formulation includes both rollout- and rule-based components, each weighted equally with  $\lambda=1$  and  $\lambda_{\text{rule}}=1$ . The reasoning task is formulated as a binary classification problem ( $|\mathcal{O}|=2$ ), using a 1-hop ( $h=1$ ) protein neighborhood from the disease-annotated subgraph. Model outputs are evaluated using precision, recall, F1-score, Jaccard similarity, Hit@5 and Hit@10, by comparing predicted target sets  $\hat{R}_n$  against reference labels  $R_n$ . Additional details are provided in the Appendix D.2.

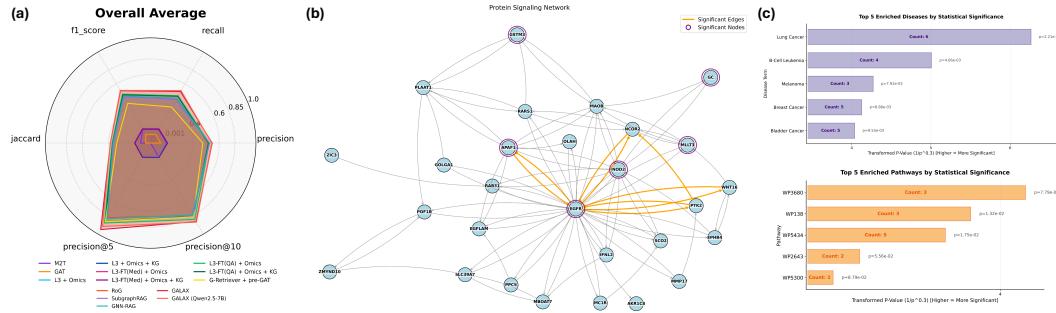
**Table 1:** Performance of models across datasets and metrics

Model	Overall		LUAD		BRCA	
	Precision $\uparrow$	Recall $\uparrow$	Precision $\uparrow$	Recall $\uparrow$	Precision $\uparrow$	Recall $\uparrow$
M2T	0.0016	0.0011	0.0020	0.0014	0.0000	0.0000
GAT	0.0006 $\pm$ 0.0000	0.0006 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0033 $\pm$ 0.0000	0.0033 $\pm$ 0.0000
L3 + Omics	0.0071 $\pm$ 0.0032	0.0013 $\pm$ 0.0002	0.0079 $\pm$ 0.0137	0.0005 $\pm$ 0.0008	0.0020 $\pm$ 0.0035	0.0017 $\pm$ 0.0029
L3 + Omics + KG	0.0125 $\pm$ 0.0032	0.0029 $\pm$ 0.0003	0.0014 $\pm$ 0.0025	0.0010 $\pm$ 0.0004	0.0073 $\pm$ 0.0068	0.0033 $\pm$ 0.0029
L3-FT(Med) + Omics	0.0179 $\pm$ 0.0045	0.0133 $\pm$ 0.0064	0.0091 $\pm$ 0.0018	0.0105 $\pm$ 0.0044	0.0110 $\pm$ 0.0088	0.0106 $\pm$ 0.0075
L3-FT(Med) + Omics + KG	0.0158 $\pm$ 0.0030	0.0058 $\pm$ 0.0011	0.0081 $\pm$ 0.0071	0.0024 $\pm$ 0.0003	0.0149 $\pm$ 0.0057	0.0050 $\pm$ 0.0000
L3-FT(QA) + Omics	0.5250 $\pm$ 0.0282	0.4959 $\pm$ 0.0435	0.5201 $\pm$ 0.0408	0.4905 $\pm$ 0.0532	0.5074 $\pm$ 0.0489	0.4856 $\pm$ 0.0570
L3-FT(QA) + Omics + KG	0.5185 $\pm$ 0.0240	0.4908 $\pm$ 0.0402	0.5214 $\pm$ 0.0242	0.4952 $\pm$ 0.0432	0.4856 $\pm$ 0.0398	0.4656 $\pm$ 0.0436
G-Retriever + pre-GAT	0.4763 $\pm$ 0.0004	0.3929 $\pm$ 0.0063	0.4642 $\pm$ 0.0181	0.3881 $\pm$ 0.0264	0.4414 $\pm$ 0.0099	0.3772 $\pm$ 0.0010
RoG	0.5248 $\pm$ 0.0134	0.4726 $\pm$ 0.0445	0.5213 $\pm$ 0.0227	0.4562 $\pm$ 0.0848	0.4791 $\pm$ 0.0375	0.4311 $\pm$ 0.0721
SubgraphRAG	0.5280 $\pm$ 0.0044	0.4617 $\pm$ 0.0027	0.5123 $\pm$ 0.0105	0.4448 $\pm$ 0.0386	0.4708 $\pm$ 0.0317	0.3917 $\pm$ 0.0376
GNN-RAG	0.5258 $\pm$ 0.0126	0.4735 $\pm$ 0.0190	0.5334 $\pm$ 0.0225	0.5052 $\pm$ 0.0170	0.4787 $\pm$ 0.0453	0.4389 $\pm$ 0.0584
<b>GALAX</b>	<b>0.5472<math>\pm</math>0.0053</b>	<b>0.5332<math>\pm</math>0.0031</b>	<b>0.5345<math>\pm</math>0.0185</b>	<b>0.5157<math>\pm</math>0.0043</b>	<b>0.5608<math>\pm</math>0.0031</b>	<b>0.5533<math>\pm</math>0.0033</b>
<b>GALAX (Qwen2.5-7B)</b>	0.5445 $\pm$ 0.0114	0.5405 $\pm$ 0.0101	0.5475 $\pm$ 0.0019	0.5462 $\pm$ 0.0111	0.5171 $\pm$ 0.0474	0.5206 $\pm$ 0.0419

**Table 2:** Hit@10 and Hit@5 for models across datasets

Model	Overall		LUAD		BRCA	
	Hit@10 $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$	Hit@5 $\uparrow$
M2T	0.0029	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0021 $\pm$ 0.0037	0.0032 $\pm$ 0.0055	0.0048 $\pm$ 0.0082	0.0095 $\pm$ 0.0165	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0122 $\pm$ 0.0033	0.0085 $\pm$ 0.0037	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0056 $\pm$ 0.0098	0.0111 $\pm$ 0.0192
L3-FT(Med) + Omics	0.0122 $\pm$ 0.0072	0.0116 $\pm$ 0.0097	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0111 $\pm$ 0.0192	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0132 $\pm$ 0.0040	0.0106 $\pm$ 0.0048	0.0048 $\pm$ 0.0082	0.0095 $\pm$ 0.0165	0.0111 $\pm$ 0.0192	0.0000 $\pm$ 0.0000
L3-FT(QA) + Omics	0.8693 $\pm$ 0.0157	0.8889 $\pm$ 0.0168	0.8667 $\pm$ 0.0218	0.8476 $\pm$ 0.0165	0.8389 $\pm$ 0.0068	<b>0.8889<math>\pm</math>0.0509</b>
L3-FT(QA) + Omics + KG	0.8529 $\pm$ 0.0153	0.8794 $\pm$ 0.0114	0.8048 $\pm$ 0.0541	0.7905 $\pm$ 0.0436	0.8222 $\pm$ 0.0347	0.8778 $\pm$ 0.0192
G-Retriever + pre-GAT	0.8550 $\pm$ 0.0046	0.8804 $\pm$ 0.0037	0.8524 $\pm$ 0.0165	0.8857 $\pm$ 0.0000	<b>0.8667<math>\pm</math>0.0000</b>	0.8667 $\pm$ 0.0000
RoG	0.8450 $\pm$ 0.0356	0.8593 $\pm$ 0.0318	0.8238 $\pm$ 0.0218	0.8095 $\pm$ 0.0436	0.7611 $\pm$ 0.1110	0.7667 $\pm$ 0.0577
SubgraphRAG	0.8476 $\pm$ 0.0167	0.8624 $\pm$ 0.0120	0.8238 $\pm$ 0.0082	0.8190 $\pm$ 0.0165	0.7333 $\pm$ 0.1014	0.7556 $\pm$ 0.0839
GNN-RAG	0.8323 $\pm$ 0.0205	0.8656 $\pm$ 0.0302	0.7571 $\pm$ 0.0623	0.7905 $\pm$ 0.0719	0.8222 $\pm$ 0.0674	0.8444 $\pm$ 0.0385
<b>GALAX</b>	0.8815 $\pm$ 0.0033	<b>0.9249<math>\pm</math>0.0048</b>	<b>0.8810<math>\pm</math>0.0082</b>	<b>0.9238<math>\pm</math>0.0436</b>	0.8500 $\pm$ 0.0441	<b>0.8889<math>\pm</math>0.0839</b>
<b>GALAX (Qwen2.5-7B)</b>	<b>0.8841<math>\pm</math>0.0126</b>	0.9079 $\pm$ 0.0084	0.8667 $\pm$ 0.0082	0.9048 $\pm$ 0.0165	0.8000 $\pm$ 0.0764	0.8556 $\pm$ 0.0385

**Baseline Models** Traditional method, **M2T** (Multiomic2Target(Deng et al., 2024)), serves as a baseline that uses only multi-omic features without graph or language modeling and performs poorly across all metrics. And we perform ablation studies to isolate the contribution of core components of language and graph modules in GALAX. On the language axis, a non-task-tuned LLaMA3 (**L3+Omics**) is weak; domain-adaptive finetuning on biomedical text (**L3-FT(Med)+Omics**) yields modest gains; task-adaptive finetuning on Target-QA (**L3-FT(QA)+Omics**) produces the step-change. On the graph axis, **GAT** incorporates graph foundation model and trained to predict the



**Figure 4: Model performance and analysis.** (a) Overall performance across metrics. (b) LUAD (ACH-000860): cancer-relevant subgraph highlighting disease-associated nodes (purple) and enrichment-supported edges (orange). (c) Enrichment analysis: top pathway and disease terms with  $p$ -values and gene counts.

CRISPR knockout effects, shows limited improvements. Furthermore, integrating graph modules into language models comes with different outputs. Adding a static KG to each language foundation models ([L3 / L3-FT(Med) / L3-FT(QA)]+Omics+KG) barely improve or even decrease model performances, and graph retrieval with pretrained GAT (G-Retriever+pre-GAT) outperform some task-adaptive finetuned language models but not reliably due to the difficulty of extracting relevant subgraphs from millions of nodes/edges. **RoG**, **SubgraphRAG** (Li et al., 2024), and **GNN-RAG** (Mavromatis & Karypis, 2025) augment the language model by retrieving optimal paths, achieving moderate improvements over L3-FT(QA)+Omics+KG. Reinforcement-guided subgraph construction on top of QA-tuned language (**GALAX: L3-FT(QA)+Omics+KG+RL**) delivers consistent, cross-dataset gains of roughly 2%-5% on each metric, indicating that a reinforcement-guided subgraph generator outperforms other graph augmented models across all datasets by enabling process-level reasoning over biologically plausible subgraphs (shown in Tables 1–2). Details of baseline models are provided in Appendix D.3.

**Computational Complexity** We denote the language model complexity by  $O(\kappa)$ . The retrieved KG subgraph includes  $M$  nodes, and  $\varepsilon$  denotes the graph embedding cost. When augmented with KG retrieval, L3-FT(QA)+Omics+KG introduces an additional  $O(M^2)$  per query, yielding  $O(\kappa + M^2)$  for both training and inference. G-retriever+preGAT embeds all  $M$  nodes and thus incurs  $O(M\varepsilon + M^2\varepsilon)$ . RoG and GNN-RAG follow the same cost at retrieval, since they embed entities and relations, requiring  $O(M\varepsilon + M^2\varepsilon)$ . SubgraphRAG reduces this by retrieving only relation triplets, which costs  $O(M^2\varepsilon)$  at retrieval. And G-retriever. RoG, GNN-RAG and SubgraphRAG all requires an  $O(\kappa)$  at both training and inference for language models. GALAX augments the language model with reinforcement-guided subgraph construction, with embedding cost  $O(M\varepsilon + M^2\varepsilon)$ . The model requires an  $O(\kappa)$  forward pass to initialize the top  $\eta$  candidates and another for final answer generation; since  $\eta \ll M$ , the graph-embedding term dominates the RL reward cost. Overall, the training and inference complexity of GALAX is therefore  $O(\kappa + M\varepsilon + M^2\varepsilon)$  (see Table 3).

**Main Results** Tables 1–2 summarize the performance of GALAX and several competitive baselines on the full test dataset. GALAX outperforms all baselines on every metric, reaching an overall precision of **0.5472** and recall of **0.5332**. To further assess target prioritization quality, we report Hit@10 and Hit@5 in Table 2, where GALAX again achieves the highest accuracy with an overall Hit@10 of **0.8815** and Hit@5 of **0.9249**. We also replaced the backbone with Qwen2.5–7B–Instruct and observed similar performances on all metrics, indicating that GALAX maintains stable performance under backbone changes (See Appendix D.4 for full experiment results). We further tested generalization by forming three holdout sets in which the selected TCGA cancer types were excluded from training and treated as unseen during evaluation. Across all holdout sets, GALAX showed only modest performance declines, indicating that it preserves strong accuracy and generalizes reliably to previously unseen cancer types (Details are provided in Appendix D.5). Beyond Target-QA, we evaluate GALAX on the pediatric cancer dataset from PedDep (Dharia et al., 2021), which offers multi-omic profiles and CRISPR-based targets for 31 tumor cell lines, using zero-shot inference due to its small sample size. Even under this setting, GALAX surpasses all baselines on all metrics,

Model	Training & Inference
L3-FT(QA)+Omics	$O(\kappa)$
L3-FT(QA)+Omics+KG	$O(\kappa + M^2)$
SubgraphRAG	$O(\kappa + M^2\varepsilon)$
G-retriever+preGAT	$O(\kappa + M\varepsilon + M^2\varepsilon)$
RoG	$O(\kappa + M\varepsilon + M^2\varepsilon)$
GNN-RAG	$O(\kappa + M\varepsilon + M^2\varepsilon)$
GALAX	$O(\kappa + M\varepsilon + M^2\varepsilon)$

**Table 3: Complexity comparisons**

**Table 4:** Overall performances for ablation studies on omic inputs and KG structure

Config	Setting	Recall $\uparrow$	Precision $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
<b>GALAX</b>					
E20	Drop 20% edges	0.5061 $\pm$ 0.0268	0.5362 $\pm$ 0.0059	0.9005 $\pm$ 0.0128	0.8741 $\pm$ 0.0060
E40	Drop 40% edges	0.2871 $\pm$ 0.0060	0.5079 $\pm$ 0.0039	0.8762 $\pm$ 0.0138	0.8471 $\pm$ 0.0090
E60	Drop 60% edges	0.2753 $\pm$ 0.0020	0.4961 $\pm$ 0.0059	0.8635 $\pm$ 0.0084	0.8307 $\pm$ 0.0033
E80	Drop 80% edges	0.2775 $\pm$ 0.0134	0.4943 $\pm$ 0.0177	0.8434 $\pm$ 0.0524	0.8185 $\pm$ 0.0494
N20	Drop 20% nodes	0.2697 $\pm$ 0.0013	0.5034 $\pm$ 0.0014	0.8786 $\pm$ 0.0070	0.8560 $\pm$ 0.0092
N40	Drop 40% nodes	0.2673 $\pm$ 0.0033	0.4901 $\pm$ 0.0035	0.8878 $\pm$ 0.0066	0.8503 $\pm$ 0.0048
N60	Drop 60% nodes	0.2617 $\pm$ 0.0056	0.4929 $\pm$ 0.0044	0.8698 $\pm$ 0.0290	0.8385 $\pm$ 0.0247
N80	Drop 80% nodes	0.2653 $\pm$ 0.0032	0.4825 $\pm$ 0.0098	0.8341 $\pm$ 0.0105	0.8103 $\pm$ 0.0134
Omic-M	Remove epigenomic data	0.4810 $\pm$ 0.0137	0.5163 $\pm$ 0.0088	0.8857 $\pm$ 0.0145	0.8614 $\pm$ 0.0115
Omic-G	Remove genomic data	0.3121 $\pm$ 0.0052	0.4277 $\pm$ 0.0056	0.8550 $\pm$ 0.0037	0.8402 $\pm$ 0.0033
Omic-T	Remove transcriptomic data	0.3377 $\pm$ 0.0016	0.4065 $\pm$ 0.0042	0.8720 $\pm$ 0.0037	0.8672 $\pm$ 0.0037
Omic-P	Remove proteomic data	0.3347 $\pm$ 0.0013	0.3980 $\pm$ 0.0058	0.8540 $\pm$ 0.0138	0.8466 $\pm$ 0.0040
Omic-All	Remove all omics	0.3024 $\pm$ 0.0032	0.3793 $\pm$ 0.0019	0.8237 $\pm$ 0.0066	0.7967 $\pm$ 0.0040
GALAX	Original	0.5332 $\pm$ 0.0031	0.5472 $\pm$ 0.0053	0.9249 $\pm$ 0.0048	0.8815 $\pm$ 0.0033
<b>L3-FT(QA) + Omic + KG</b>					
E20	Drop 20% edges	0.4599 $\pm$ 0.0820	0.5214 $\pm$ 0.0036	0.8587 $\pm$ 0.0157	0.8373 $\pm$ 0.0258
E40	Drop 40% edges	0.2742 $\pm$ 0.0090	0.5064 $\pm$ 0.0093	0.8429 $\pm$ 0.0173	0.8226 $\pm$ 0.0169
E60	Drop 60% edges	0.2676 $\pm$ 0.0062	0.4991 $\pm$ 0.0099	0.8296 $\pm$ 0.0305	0.8111 $\pm$ 0.0175
E80	Drop 80% edges	0.2611 $\pm$ 0.0074	0.4880 $\pm$ 0.0086	0.8254 $\pm$ 0.0361	0.8063 $\pm$ 0.0370
N20	Drop 20% nodes	0.2662 $\pm$ 0.0059	0.4916 $\pm$ 0.0024	0.8434 $\pm$ 0.0186	0.8222 $\pm$ 0.0193
N40	Drop 40% nodes	0.2658 $\pm$ 0.0049	0.4838 $\pm$ 0.0161	0.8709 $\pm$ 0.0422	0.8339 $\pm$ 0.0335
N60	Drop 60% nodes	0.2648 $\pm$ 0.0000	0.4742 $\pm$ 0.0181	0.7857 $\pm$ 0.0247	0.7690 $\pm$ 0.0303
N80	Drop 80% nodes	0.2689 $\pm$ 0.0042	0.4722 $\pm$ 0.0065	0.7111 $\pm$ 0.0055	0.6974 $\pm$ 0.0111
Omic-M	Remove epigenomic data	0.4602 $\pm$ 0.0361	0.4878 $\pm$ 0.0256	0.8794 $\pm$ 0.0055	0.8466 $\pm$ 0.0142
Omic-G	Remove genomic data	0.3213 $\pm$ 0.0047	0.3962 $\pm$ 0.0098	0.8455 $\pm$ 0.0073	0.8349 $\pm$ 0.0136
Omic-T	Remove transcriptomic data	0.3244 $\pm$ 0.0034	0.3996 $\pm$ 0.0087	0.8550 $\pm$ 0.0073	0.8381 $\pm$ 0.0097
Omic-P	Remove proteomic data	0.3266 $\pm$ 0.0012	0.3872 $\pm$ 0.0032	0.8497 $\pm$ 0.0048	0.8265 $\pm$ 0.0056
Omic-All	Remove all omics	0.2669 $\pm$ 0.0075	0.3577 $\pm$ 0.0098	0.7830 $\pm$ 0.0182	0.7538 $\pm$ 0.0174
L3-FT(QA)+Omic+KG	Original	0.4908 $\pm$ 0.0402	0.5185 $\pm$ 0.0240	0.8794 $\pm$ 0.0114	0.8529 $\pm$ 0.0153

showing strong transfer to external dataset (Details are provided in Appendix D.6). Figure 4b-c illustrates the explainable subgraph generated for the lung cancer cell line ACH-000860. To further validate the biological relevance of the extracted subgraph, we performed functional enrichment analysis. The results reveal significant enrichment in cancer-associated signaling pathways, including the cancer pathway WP5434 and EGFR-related receptor signaling pathways such as WP138 and WP3680, as cataloged in WikiPathways (Agrawal et al., 2024). Notably, EGFR a well-established therapeutic target in NSCLC (Steuer & Ramalingam, 2015) appears in five enriched terms, together with PTK2 and WNT16 which are known to regulate invasion, epithelial mesenchymal transition, and therapeutic resistance (Tong et al., 2019; Sun et al., 2012). The observed pathway enrichment provides strong biological support for the relevance of the selected targets in lung cancer. Additional disease enrichment using the GAD DISEASE database further supports this conclusion (Sherman et al., 2022), with lung cancer identified as the top associated disease term with a p value of 0.0022, involving GSTM3, APAF1, NOD2, MLLT3, GC and EGFR. Full enrichment details across all cancer types are included in the Appendix E.1. In addition, human and LLM evaluation results in Appendix E.2 show that most generated subgraphs are biologically plausible.

**Ablation Studies** As shown in Table 4, we evaluated GALAX under systematic perturbations to omic inputs and KG structure. The KG provides both the guidance for graph generation and the retrieved subgraphs that supply biological context to the language model, so removing portions of the KG naturally reduces overall performance. Under KG deletions, GALAX remained stronger than the baseline across all edge-removal levels because node attributes and pretrained GNN embeddings enabled the model to assemble meaningful subnetworks from sparse structure. Node deletion was more destructive since removing 20% of nodes removed about 35% of edges, but GALAX maintained more stable Hit@5 and Hit@10 scores than the baseline across all deletion levels. Removing epigenomic data produced the smallest decline due to methylation saturation, while removing genomic, transcriptomic, or proteomic signals caused larger drops, yet GALAX still outperformed the L3-FT(QA)+Omic+KG baseline even when all omics were removed. Overall, these findings show that GALAX generalizes well under shifts in omic distributions and reduced KG connectivity.

## 6 CONCLUSION

We present **GALAX**, a graph-augmented language model that unifies numerical multi-omic evidence, literature-scale textual information, and stepwise graph construction under reinforcement learning with biologically grounded supervision through a Graph Process Reward Model (GRPM), which scores intermediate subgraphs for biological plausibility and cancer relevance and guides the system to generate patient-specific mechanistic subgraph for target prioritization without explicit labels. To facilitate evaluation, we introduce **Target-QA**, a benchmark combining multi-omic data, CRISPR outcomes, and graph knowledge for target discovery. GALAX consistently outperforms baseline models on this dataset.

## ETHICS STATEMENT

This dataset is derived from the Broad Institute’s Cancer Dependency Map (DepMap<sup>1</sup>) and will be released strictly for non-commercial, internal research and academic use, consistent with DepMap’s Terms of Use. We do not redistribute original DepMap files; instead, we provide derived, non-identifiable annotations and processing scripts/pointers so users can obtain the source data directly from DepMap after accepting its terms. The dataset is not intended for clinical applications and must not be used for any Commercial Use (e.g., direct sale, incorporation into a product, or training/developing/enhancing ML/AI models beyond internal academic research). Users agree to acknowledge DepMap and the Broad Institute using the acknowledgement wording specified by DepMap, and to respect any third-party rights that may attach to the underlying data. Users must preserve confidentiality and refrain from any re-identification attempts. This statement summarizes our compliance posture and does not constitute legal advice; users are responsible for ensuring their own compliance with DepMap’s Terms and applicable policies.

## REPRODUCIBILITY STATEMENT

We release the source code together with preprocessing pipelines, thereby enabling reproduction of Target-QA datasets and reported experiments. The **Target-QA**<sup>2</sup> and **GALAX**<sup>3</sup> are publicly available at Huggingface and GitHub.

## ACKNOWLEDGMENTS

This research was partially supported by NLM 1R01LM013902-01A1.

## REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Ayushi Agrawal, Hasan Balcı, Kristina Hanspers, Susan L Coort, Marvin Martens, Denise N Slenter, Friederike Ehrhart, Daniela Digles, Andra Waagmeester, Isabel Wassink, et al. Wikipathways 2024: next generation pathway database. *Nucleic acids research*, 52(D1):D679–D689, 2024.

Ricard Argelaguet, Britta Velten, Damien Arnol, Sascha Dietrich, Thorsten Zenz, John C Marioni, Florian Buettner, Wolfgang Huber, and Oliver Stegle. Multi-omics factor analysis—a framework for unsupervised integration of multi-omics data sets. *Molecular systems biology*, 14(6):e8124, 2018.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

Francis Crick. Central dogma of molecular biology. *Nature*, 227(5258):561–563, 1970.

Joshua M Dempster, Jordan Rossen, Mariya Kazachkova, Joshua Pan, Guillaume Kugener, David E Root, and Aviad Tsherniak. Extracting biological insights from the project achilles genome-scale crispr screens in cancer cell lines. *BioRxiv*, pp. 720243, 2019.

Eden Z Deng, Giacomo B Marino, Daniel JB Clarke, Ido Diamant, Adam C Resnick, Weiping Ma, Pei Wang, and Avi Ma’ayan. Multiomics2targets identifies targets from cancer cohorts profiled with transcriptomics, proteomics, and phosphoproteomics. *Cell Reports Methods*, 4(8), 2024.

<sup>1</sup><https://depmap.org/portal/>

<sup>2</sup>Hugging Face: <https://huggingface.co/datasets/FuhaiLiAiLab/Target-QA>

<sup>3</sup>GitHub: <https://github.com/FuhaiLiAiLab/GALAX>

Neekesh V Dharia, Guillaume Kugener, Lillian M Guenther, Clare F Malone, Adam D Durbin, Andrew L Hong, Thomas P Howard, Pratiti Bandopadhayay, Caroline S Wechsler, Iris Fung, et al. A first-generation pediatric cancer dependency map. *Nature genetics*, 53(4):529–538, 2021.

Bahare Fatemi, Jonathan Halcrow, and Bryan Perozzi. Talk like a graph: Encoding graphs for large language models. *arXiv preprint arXiv:2310.04560*, 2023.

Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, pp. 10835–10866. PMLR, 2023.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Yehudit Hasin, Marcus Seldin, and Aldons Lusis. Multi-omics approaches to disease. *Genome biology*, 18(1):83, 2017.

Xiaoxin He, Yijun Tian, Yifei Sun, Nitesh Chawla, Thomas Laurent, Yann LeCun, Xavier Bresson, and Bryan Hooi. G-retriever: Retrieval-augmented generation for textual graph understanding and question answering. *Advances in Neural Information Processing Systems*, 37:132876–132907, 2024.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Tomasz Korbak, Kejian Shi, Angelica Chen, Rasika Vinayak Bhalerao, Christopher Buckley, Jason Phang, Samuel R Bowman, and Ethan Perez. Pretraining language models with human preferences. In *International Conference on Machine Learning*, pp. 17506–17533. PMLR, 2023.

Vessela N Kristensen, Ole Christian Lingjærde, Hege G Russnes, Hans Kristian M Vollan, Arnoldo Frigessi, and Anne-Lise Børresen-Dale. Principles and methods of integrative genomic analyses in cancer. *Nature Reviews Cancer*, 14(5):299–313, 2014.

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.

Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474, 2020.

Mufei Li, Siqi Miao, and Pan Li. Simple is effective: The roles of graphs and large language models in knowledge-graph-based retrieval-augmented generation. *arXiv preprint arXiv:2410.20724*, 2024.

Tianxiang Li, Yanyan Yang, Hongzhao Qi, Weigang Cui, Lin Zhang, Xiuxiu Fu, Xiangqin He, Meixin Liu, Pei-feng Li, and Tao Yu. Crispr/cas9 therapeutics: progress and prospects. *Signal transduction and targeted therapy*, 8(1):36, 2023.

Xiao Li, Jie Ma, Ling Leng, Mingfei Han, Mansheng Li, Fuchu He, and Yunping Zhu. Mogen: a multi-omics integration method based on graph convolutional network for cancer subtype analysis. *Frontiers in Genetics*, 13:806842, 2022.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.

Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Meiqi Guo, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, Lei Meng, et al. Improve mathematical reasoning in language models by automated process supervision. *arXiv preprint arXiv:2406.06592*, 2024.

Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. Reasoning on graphs: Faithful and interpretable large language model reasoning. *arXiv preprint arXiv:2310.01061*, 2023.

Costas Mavromatis and George Karypis. Gnn-rag: Graph neural retrieval for efficient large language model reasoning on knowledge graphs. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 16682–16699, 2025.

Chen Meng, Oana A Zelezniak, Gerhard G Thallinger, Bernhard Kuster, Amin M Gholami, and Aedin C Culhane. Dimension reduction techniques for the integrative analysis of multi-omics data. *Briefings in bioinformatics*, 17(4):628–641, 2016.

Nam D Nguyen and Daifeng Wang. Multiview learning for understanding functional multiomics. *PLoS computational biology*, 16(4):e1007677, 2020.

Xubin Ren, Jiabin Tang, Dawei Yin, Nitesh Chawla, and Chao Huang. A survey of large language models for graphs. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6616–6626, 2024.

Florian Rohart, Benoît Gautier, Amrit Singh, and Kim-Anh Lê Cao. mixomics: An r package for ‘omics feature selection and multiple data integration. *PLoS computational biology*, 13(11):e1005752, 2017.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Ophir Shalem, Neville E Sanjana, Ella Hartenian, Xi Shi, David A Scott, Tarjei S Mikkelsen, Dirk Heckl, Benjamin L Ebert, David E Root, John G Doench, et al. Genome-scale crispr-cas9 knockout screening in human cells. *Science*, 343(6166):84–87, 2014.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

Ronglai Shen, Adam B Olshen, and Marc Ladanyi. Integrative clustering of multiple genomic data types using a joint latent variable model with application to breast and lung cancer subtype analysis. *Bioinformatics*, 25(22):2906–2912, 2009.

Brad T Sherman, Ming Hao, Ju Qiu, Xiaoli Jiao, Michael W Baseler, H Clifford Lane, Tomozumi Imamichi, and Weizhong Chang. David: a web server for functional enrichment analysis and functional annotation of gene lists (2021 update). *Nucleic acids research*, 50(W1):W216–W221, 2022.

Junwei Shi, Eric Wang, Joseph P Milazzo, Zihua Wang, Justin B Kinney, and Christopher R Vakoc. Discovery of cancer drug targets by crispr-cas9 screening of protein domains. *Nature biotechnology*, 33(6):661–667, 2015.

Conor E Steuer and Suresh S Ramalingam. Targeting egfr in lung cancer: Lessons learned and future perspectives. *Molecular aspects of medicine*, 45:67–73, 2015.

Sandra Steyaert, Marija Pizurica, Divya Nagaraj, Priya Khandelwal, Tina Hernandez-Boussard, Andrew J Gentles, and Olivier Gevaert. Multimodal data fusion for cancer biomarker discovery with deep learning. *Nature machine intelligence*, 5(4):351–362, 2023.

Yu Sun, Judith Campisi, Celestia Higano, Tomasz M Beer, Peggy Porter, Ilsa Coleman, Lawrence True, and Peter S Nelson. Treatment-induced damage to the tumor microenvironment promotes prostate cancer therapy resistance through wnt16b. *Nature medicine*, 18(9):1359–1368, 2012.

Xuexia Tong, Ryosuke Tanino, Rong Sun, Yukari Tsubata, Tamio Okimoto, Mayumi Takechi, and Takeshi Isobe. Protein tyrosine kinase 2: a novel therapeutic target to overcome acquired egfr-ki resistance in non-small cell lung cancer. *Respiratory Research*, 20(1):270, 2019.

Eric J Topol. High-performance medicine: the convergence of human and artificial intelligence. *Nature medicine*, 25(1):44–56, 2019.

Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models. *arXiv preprint arXiv:2305.04091*, 2023.

Tongxin Wang, Wei Shao, Zhi Huang, Haixu Tang, Jie Zhang, Zhengming Ding, and Kun Huang. Mogenet integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification. *Nature communications*, 12(1):3445, 2021.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M Stuart. The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10):1113–1120, 2013.

Heming Zhang, Yixin Chen, Philip Payne, and Fuhai Li. Using deepsignalingflow to mine signaling flows interpreting mechanism of synergy of cocktails. *npj Systems Biology and Applications*, 10(1):92, 2024a.

Heming Zhang, Shunning Liang, Tim Xu, Wenyu Li, Di Huang, Yuhang Dong, Guangfu Li, J Philip Miller, S Peter Goedegebuure, Marco Sardiello, et al. Biomedgraphica: An all-in-one platform for biomedical prior knowledge and omic signaling graph generation. *bioRxiv*, 2024b.

Heming Zhang, Tim Xu, Dekang Cao, Shunning Liang, Lars Schimmelpfennig, Levi Kaster, Di Huang, Carlos Cruchaga, Guangfu Li, Michael Province, et al. Omnicelltoscg: The first cell text-omic signaling graphs dataset for joint llm and gnn modeling. *arXiv preprint arXiv:2504.02148*, 2025a.

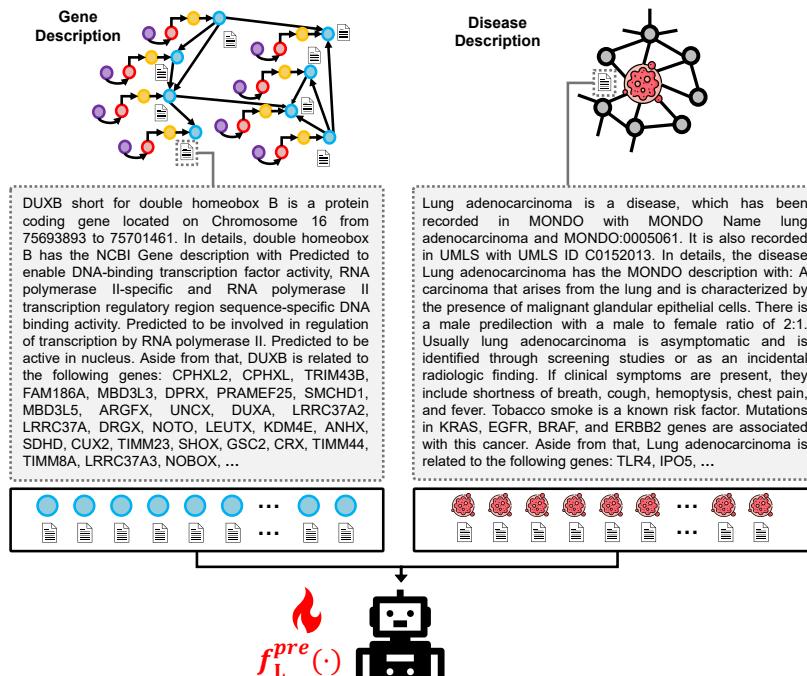
Jingyi Zhang, Jiaxing Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv preprint arXiv:2503.12937*, 2025b.

## A THE USE OF LARGE LANGUAGE MODELS

We used ChatGPT-5 as a writing assistant. All LLM-suggested text was reviewed, fact-checked, and edited by the authors, who take full responsibility for the final content. The LLM is not an author and is not eligible for authorship under the ICLR Code of Ethics.

## B DATASET

### B.1 BIOMEDICAL TERMINOLOGY CORPUS



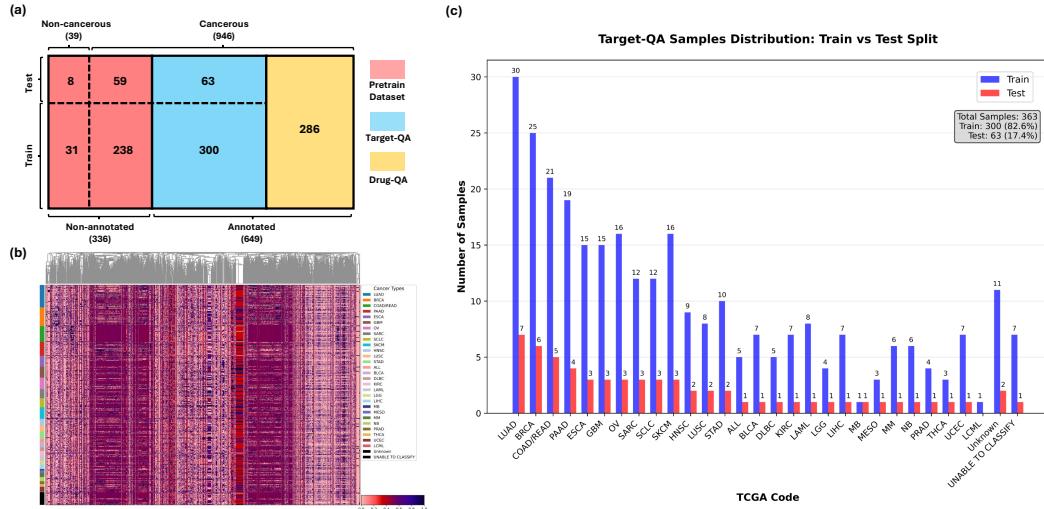
**Figure 5: Biomedical Corpus Composition.** Pretraining LLM  $f_L^{\text{pre}}(\cdot)$  on biomedical terminology and structural patterns using protein-protein and disease-protein interactions.

Each entities or nodes in TOSG has the name and description about it with  $\mathcal{T} = \{T_{\text{name}}, T_{\text{desc}}\}$ , where  $|T_{\text{name}}| = |T_{\text{desc}}| = M$ . And disease entity has has an associated name and textual description with  $\mathcal{S} = \{S_{\text{name}}, S_{\text{desc}}\}$  where  $|\mathcal{S}| = M^{(S)}$ . In the Figure 5, we pretrained  $f_L^{\text{pre}}$  with curated text copora in BioMedGraphica<sup>4</sup>, where they provided the data collection and integration source code. Following their processed descriptions, we incorporate  $\mathcal{G}^{(\text{PPI})}$  and  $\mathcal{G}^{(\text{DTI})}$  to enrich the corpus by appending protein-protein (PPI) and disease-protein (DTI) interaction information as textual descriptions after each protein and disease entity. For entities without known interactions, we assign empty strings during corpus construction, resulting in the intermediate representations  $\mathcal{T}$  and  $\mathcal{S}$ . In practice, we exclude these empty entities to derive the final input sets  $\mathcal{T}'$  and  $\mathcal{S}'$ , where  $\mathcal{T}'$  includes 42,224 protein descriptions and  $\mathcal{S}'$  includes 22,340 disease descriptions (i.e.,  $|\mathcal{T}'| = 42,224$  and  $|\mathcal{S}'| = M'^{(S)} = 22,340$ ). This yields a combined corpus of 64,564 text samples for pretraining.

### B.2 DEPMAP DATA PREPROCESSING

As shown in Figure 6a, after multi-omics integration in BioMedGraphica of DepMap cohort (see Table 6) comprises  $N^{(0)}$  ( $N^{(0)} = 985$ ) cell-line samples, which are organized into three datasets: a pretraining set, Target-QA, and Drug-QA. Of these 985 samples, 336 samples lack disease/tcga-code annotations or belong to non-cancerous samples, while 649 samples cancerous. Within the annotated cancer set, 363 samples overlap with DepMap CRISPR multi-omic data. Since we would like to utilize as many as samples for training GALAX, we set  $N = 363$  as Target-QA,  $N^{(0)} = 336$

<sup>4</sup><https://huggingface.co/datasets/FuhaiLiAiLab/BioMedGraphica>



**Figure 6: Dataset composition, feature landscape, and split statistics.** (a) Block diagram of the corpus partitioned by phenotype (non-cancerous vs. cancerous) and by task: pretraining pool (pink), Target-QA (blue), and Drug-QA (yellow). Numbers within boxes denote sample counts; the dashed horizontal line marks the train/test division and the dashed vertical line marks the non-cancerous/cancerous and non-annotated/annotated boundaries. (b) Heatmap of standardized multi-omic features (top 5,000 most variance features) for Target-QA samples. Rows are samples (annotated by TCGA cancer type); columns are features with top variances. The colored sidebar encodes TCGA cancer types (legend at right). (c) Distribution of Target-QA samples by TCGA code for train (blue) and test (red); totals shown in the inset ( $N = 363$ ; train = 300 (82.6%), test = 63 (17.4%)).

**Table 5:** TCGA cancer type codes and their full names.

TCGA Code	Full Name	TCGA Code	Full Name
Overall Average	Overall Average	MB	Medulloblastoma
LUAD	Lung Adenocarcinoma	ALL	Acute Lymphoblastic Leukemia
BRCA	Breast Invasive Carcinoma	LGG	Brain Lower Grade Glioma
COAD/READ	Colon/Rectum Adenocarcinoma	NB	Neuroblastoma
PAAD	Pancreatic Adenocarcinoma	MESO	Mesothelioma
GBM	Glioblastoma Multiforme	LIHC	Liver Hepatocellular Carcinoma
SARC	Sarcoma	LAML	Acute Myeloid Leukemia
OV	Ovarian Serous Cystadenocarcinoma	DLBC	Lymphoid Neoplasm Diffuse Large B-cell Lymphoma
SKCM	Skin Cutaneous Melanoma	MM	Multiple Myeloma
ESCA	Esophageal Carcinoma	KIRC	Kidney Renal Clear Cell Carcinoma
SCLC	Small Cell Lung Cancer	THCA	Thyroid Carcinoma
HNSC	Head and Neck Squamous Cell Carcinoma	BLCA	Bladder Urothelial Carcinoma
LUSC	Lung Squamous Cell Carcinoma	UCEC	Uterine Corpus Endometrial Carcinoma
STAD	Stomach Adenocarcinoma	PRAD	Prostate Adenocarcinoma

as pretraining set. To pretrain the graph foundation model,  $f_G^{\text{pre}}, f_G$ , we consider  $|\mathcal{O}| = 2$  classes—cancerous (297 samples) and non-cancerous (39 samples)—and perform an 80/20 random split, yielding 269 training samples (238 cancerous, 31 non-cancerous) and 67 test samples (59 cancerous, 8 non-cancerous). The raw omics feature matrices include promoter and gene modalities  $m^{(pm)}$  and  $m^{(g)}$  (each with 86,238 entities), a transcript modality  $m^{(t)}$  (412,039 transcript-level entities), and a protein modality  $m^{(p)}$  (121,419 protein-level entities). The promoter modality is represented as a virtual node type in the graph encoder—entity-wise duplicates of genes whose omic values are drawn from DepMap methylation. Together, these modalities yield  $M = 834,809$  omics entities (see Table 7). The unified knowledge graph integrates protein–protein interactions ( $|\mathcal{E}| = 27,087,971$ ) and disease–target associations ( $|\mathcal{E}^{(\text{PPI})}| = 17,151,453$ ). Every node in the Text-Omic Signaling Graph (TOSG) carries text attributes  $\mathcal{T} = \{T_{\text{name}}, T_{\text{desc}}\}$  with  $|T_{\text{name}}| = |T_{\text{desc}}| = M$ ; empty fields are set to the empty string to preserve schema alignment. Hence, we construct the TOSG by linking multi-omics features with biomedical relational knowledge using BioMedGraphica and pretrain the graph encoder  $f_G^{\text{pre}}, f_G$  (parameters  $\theta_G^{\text{pre}}, \theta_G$ ) with a masked-edge modeling objective that encourages recovery of held-out interactions from context, capturing implicit omic relationships and signaling dependencies.

**Table 6:** Descriptions and sources of raw data files used in this study

File Type	File Name	Download Site
Promoter feature	CCLE_RRBS_TSS1kb_20181022.txt	<a href="https://depmap.org/portal/data_page/?tab=allData">https://depmap.org/portal/data_page/?tab=allData</a>
Gene feature	OmicsCNGene.csv	<a href="https://depmap.org/portal/data_page/?tab=allData">https://depmap.org/portal/data_page/?tab=allData</a>
Transcript feature	See footnote <sup>†</sup>	<a href="https://depmap.org/portal/data_page/?tab=allData">https://depmap.org/portal/data_page/?tab=allData</a>
Protein feature	protein.quant.current.normalized.csv	<a href="https://depmap.org/portal/data_page/?tab=allData">https://depmap.org/portal/data_page/?tab=allData</a>
CRISPR gene effect	CRISPRGeneEffect.csv	<a href="https://depmap.org/portal/data_page/?tab=allData">https://depmap.org/portal/data_page/?tab=allData</a>
Cell line annotation	Table_S1_Sample_Information.xlsx	<a href="https://depmap.org/portal/data_page/?tab=allData">https://depmap.org/portal/data_page/?tab=allData</a>
Cell line annotation	cellosaurus.obo	<a href="https://ftp.expasy.org/databases/cellosaurus/cellosaurus.obo">https://ftp.expasy.org/databases/cellosaurus/cellosaurus.obo</a>
Cell line status	cell-lines-in-Non-Cancerous.csv	<a href="https://depmap.org/portal/context/Non-Cancerous">https://depmap.org/portal/context/Non-Cancerous</a>

<sup>†</sup> Transcript file name: OmicsExpressionProteinCodingGenesTPMLog1BatchCorrected.csv

**Table 7:** Summary of feature dimensions across omics and samples

Modality	Raw Matrix	Processed Matrix
Promoter	21,337 rows, 846 samples	86,238 entities, 985 samples
Gene	38,590 rows, 1,928 samples	86,238 entities, 985 samples
Transcript	19,138 rows, 1,672 samples	412,039 entities, 985 samples
Protein	12,755 rows, 378 samples	121,419 entities, 985 samples
Cell Line Annotation	1,019 samples	985 samples
Non-cancer Samples	137 samples	39 samples

### B.3 TARGET-QA GENERATION

Based on the raw data provided from the DepMap CRISPR gene effect data with 1178 samples, we get the overlapped 363 ( $N = 363$ ) samples with the pretraining samples. And we do 80/20 train/test split with 80/20 ratio for 300 training samples and 63 test samples at random seeds. In total, we collected test samples from LUAD (7 samples), BRCA (6 samples), COAD/READ (5 samples), PAAD (4 samples), ESCA (3 samples), GBM (3 samples), OV (3 samples), SARC (3 samples), SCLC (3 samples), SKCM (3 samples), HNSC (2 samples), LUSC (2 samples), STAD (2 samples), etc. (see Figure 6C). Given that methylation values in DepMap have so many are 1 (full methylated over the promoter region around the transcription start site), which means that many values are ranked as top  $K$  ( $K=10$ ), so we just omit the methylation (epigenomic) values by only providing genomic, transcriptomic and proteomic values. Afterwards, each QA sample is indexed by a unique key corresponding to the cancer cell line, such as:

- ACH-000098: The identifier for a glioblastoma cancer cell line.

The corresponding JSON object contains the following fields:

- **cell\_line\_name**: Name of the cancer cell line (e.g., GAMG).
- **sample\_dti\_index**: Index for omics numpy data to be fetched.

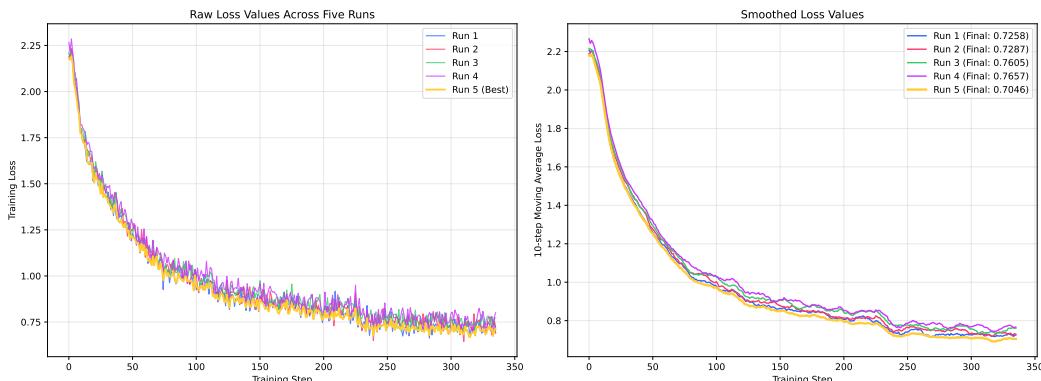
- **disease**: Name of the associated disease (e.g., glioblastoma).
- **disease\_bmgc\_id**: BioMedGraphica-Conn identifier of the disease (e.g., BMGC\_DS00965).
- **input**:
  - **top\_k\_gene, top\_k\_transcript, top\_k\_protein**:
    - \* `hgnc_symbols`: List of gene/transcript/protein names.
    - \* `protein_bmgc_ids`: Corresponding BioMedGraphica-Conn identifiers.
    - \* `protein_llmname_ids`: Other synonymy names or IDs for corresponding genes/transcripts/proteins.
  - **knowledge\_graph**:
    - \* `disease_protein`: Includes `bmgc_ids`, `hgnc_symbols`, and indices for disease-associated proteins.
    - \* `ppi_neighbors`: PPI-linked proteins with similar structure as above.
    - \* `protein_relationships`: Textual descriptions of biological interactions (e.g., "BRCA1 → TP53").
- **ground\_truth\_answer**: Contains the validated target(s) used for evaluation:
  - `hgnc_symbols`: HGCN symbol names for CRISPR targets
  - `protein_bmgc_ids`: Corresponding CRISPR targets names for BioMedGraphica-Conn names
  - `protein_llmname_ids`: Other synonymy names or IDs

This hierarchical structure supports multi-modal reasoning by organizing omic features, biomedical knowledge graphs, and ground-truth target labels, where the cell line name is denoted as  $c_n$ , the disease name as  $s'_n$ , the top-ranked genes, transcripts, and proteins as  $X_n^{(K)}$ , and the associated knowledge graph as  $\mathcal{G}_n^{(\text{sub})}$ .

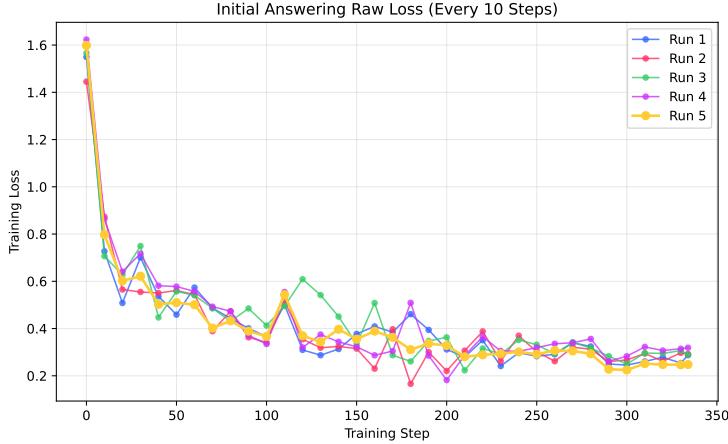
## C PRETRAINING OF FOUNDATION MODELS

### C.1 LLM PRETRAINING ON BIOMEDICAL CORPUS

We pretrain a large language model, denoted as  $f_L^{\text{pre}}$  with parameters  $\theta_L^{\text{pre}}$ , using curated text corpora from final input sets  $T'$  and  $S'$ . Hence, by pretraining the Llama3-8B-Instruct for 3 epochs using a per-device batch size of 16 and a gradient accumulation step of 8, resulting in an effective batch size of 128. The optimizer was AdamW with a learning rate of 1e-5, no weight decay, and a cosine learning rate schedule with a warm-up ratio of 10%. Gradient clipping was applied with a maximum gradient norm of 1.0. To improve memory efficiency during training, we enabled gradient checkpointing and utilized `bf16` precision while disabling `fp16`. Training proceeded for 336 steps, with the loss decreasing from approximately 2.2 at the start to around 0.7 by the end (see Figure 7).



**Figure 7:** Pretraining language model loss on biomedical corpus

**Figure 8:** Initial language answering training loss

## C.2 LLM PRETRAINING ON TARGET-QA

We will supervise the language model using the refined answer  $A_n$  as the ground-truth label from Target-QA with:

$$\mathcal{L}_{\text{init}} = - \sum_{n=1}^N \sum_{j=1}^J \log \xi_{\theta_{\text{init}}, \theta_L}(a_{n,j} \mid a_{n,<j}, P_n^{(\text{init})}) \quad (15)$$

, where  $P_n^{(\text{init})}$  is the structured prompt as shown in Table 8. This objective enables the model to refine its initial answering based on user-specific and naive graph information. In details, for each sample  $c_n$ , a question is constructed as  $Q_n = \{c_n, s'_n, X_n^{(K)}, \mathcal{G}_n^{(\text{sub})}\}$ , which integrates the cell line identifier, disease label, multi-omics features, and a knowledge subgraph specific to the sample context. The corresponding answer  $A_n$  is a sentence that enumerates the top  $\gamma$  CRISPR-prioritized gene targets for sample  $n$ , denoted as  $R_n = \{r_{n,1}, r_{n,2}, \dots, r_{n,\gamma}\}$ . Each data instance is represented as a tuple  $D_n = (Q_n, A_n)$ , and the full dataset is given by  $\mathcal{D} = \{D_1, D_2, \dots, D_N\}$ . The input to the model is a structured prompt  $P_n^{(\text{init})}$  derived from  $Q_n$ , in which the knowledge subgraph  $\mathcal{G}_n^{(\text{sub})}$  is translated into a natural language format designed for expert-level graph reasoning. For example, a subgraph  $\mathcal{G}_n^{(\text{sub})}$  may contain nodes BRCA1, TP53, EGFR, MAPK1, AKT1, PIK3CA, MTOR, PTEN, and CDK2, with observed interactions: BRCA1  $\rightarrow$  TP53, TP53  $\rightarrow$  EGFR, EGFR  $\rightarrow$  MAPK1, EGFR  $\rightarrow$  AKT1, AKT1  $\rightarrow$  MTOR, PIK3CA  $\rightarrow$  AKT1, PIK3CA  $\rightarrow$  PTEN, PTEN  $\rightarrow$  MTOR, and MAPK1  $\rightarrow$  CDK2. This subgraph defines the molecular context for reasoning about gene knockout effects in the cell line  $c_n$  under disease condition  $s'_n$ . We conduct five independent finetuning trials using the constructed Target-QA dataset, each initialized with a different random seed. The resulting training loss trajectories are shown in Figure 8. Among these, the best-performing run (run 5) exhibits stable convergence, beginning with an initial loss of approximately 1.3 and reaching a final loss of around 0.3. We select this run as the final model checkpoint and designate it as our initial model  $f_{\text{init}}$ , which serves as the foundation for downstream reasoning and refined target prioritization.

For pretraining  $f_{\text{init}}$  with parameterized  $\theta_{\text{init}}$ , we pretrain for 5 epochs with per-device batch size 1 and gradient accumulation 2 (effective global batch size =  $2 \times$  GPUs), using AdamW (adamw\_torch) with learning rate  $1 \times 10^{-5}$ , cosine schedule with 10% warmup, and gradient clipping at 0.5. We enable gradient checkpointing and bfloat16 (with fp16 disabled), run on 2  $\times$  NVIDIA H100 80GB with DeepSpeed ZeRO-3. Checkpointing and evaluation occur every 67 steps with up to 5 checkpoints retained. Data loading uses 4 workers, pinned memory, no last-batch drop, and no length grouping; we set seed= 42.

## C.3 GRAPH FOUNDATION MODELS

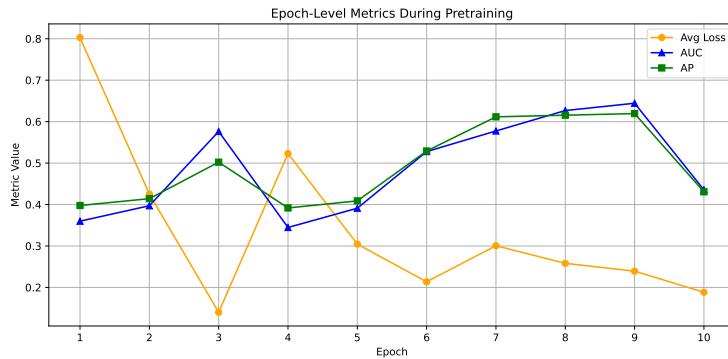
**Pretraining for Capturing the Edge Mechanism** We pretrain a graph model, denoted as  $f_G^{\text{pre}}$  with parameters  $\theta_G^{\text{pre}}$ , using pretraining samples from  $\mathcal{X}^{(0)}$ . To effectively model graph-structured biological relationships, we pretrain our model using a masked edge prediction objective com-

**Table 8:** Prompt design for  $P_n^{(\text{init})}$  and expected output for initial protein target reasoning.

Section	Content
<b>Instruction</b>	Identify the 100 priority genes whose knockout causes the strongest negative effect on the viability or proliferation of cell line $c_n$ in the context of disease $s'_n$ , based on multi-omics and knowledge graph signals.
<b>Input</b>	<ul style="list-style-type: none"> <li>- Top 10 ranked genes with amplification from copy number data: <math>g_1'^{(n)}, g_2'^{(n)}, \dots, g_K'^{(n)}</math></li> <li>- Top 10 ranked transcripts with high expression: <math>t_1'^{(n)}, t_2'^{(n)}, \dots, t_K'^{(n)}</math></li> <li>- Top 10 ranked proteins from RPPA: <math>p_1'^{(n)}, p_2'^{(n)}, \dots, p_K'^{(n)}</math></li> <li>- Disease-associated proteins from the knowledge graph: <math>\mathcal{V}_n^{(\text{sub})}</math></li> <li>- Known protein–protein/disease–protein relationships: <math>\mathcal{E}_n^{(\text{sub})}</math></li> </ul>
<b>Output</b>	<p>Based on the integrated multi-omics data and knowledge graph, I identified the 100 genes whose knockout is predicted to have the most severe negative impact on the viability or proliferation of the <math>c_n</math> cell line in <math>s'_n</math>. The prioritized gene list is as follows:</p> <ol style="list-style-type: none"> <li>1. <math>r_{n,1}^{(\text{init})}</math></li> <li>2. <math>r_{n,2}^{(\text{init})}</math></li> <li>3. <math>r_{n,3}^{(\text{init})}</math></li> <li>...</li> </ol> <p>These genes represent critical vulnerabilities for the given cell line under the disease context.</p>

bined with random walk-based graph sampling. The process starts by masking a small fraction ( $p = 0.0001$ ) of edges, allowing the model to infer missing interactions based on surrounding omics and text-derived features. We adopt a GAT-based encoder with two layers, each consisting of 8 hidden channels. Decoder layers use 4 channels. Both encoder and decoder modules apply dropout at a rate of 0.2. The model supports optional batch normalization and uses a `leaky_relu` activation function. An internal encoder stack of up to 4 layers enables deeper relational modeling. All architecture and training options are managed through a reproducible `argparse` interface. Multimodal inputs include one omic feature and a text embedding of dimension 1, initialized using BioBERT v1.1. The model optionally supports training of the text encoder via `train_text`. We use a pretraining batch size of 4 for omics data and 64 for text. The optimizer is AdamW with a learning rate of 0.001, weight decay of  $5 \times 10^{-5}$ , and gradient norm clipping at 1.0.

The pretrained model is evaluated using average loss, AUC, and average precision (AP) over validation batches. As shown in Figure 9, the model demonstrates progressive improvement across steps, reaching a minimum batch loss of 0.140, peak AUC of 0.644, and peak AP of 0.619. These results confirm that the model successfully learns meaningful edge semantics and multimodal associations during pretraining. All training is conducted on a single NVIDIA H100 GPU (80GB).

**Figure 9:** Pretraining loss to capture gene regulatory mechanism

**Pretraining for Capturing the Cancerous Status** Due to the severe class imbalance in the training set, we apply random oversampling to the minority class (non-cancerous) to balance the class

**Table 9:** Refined prompt design for  $P_n^{(\text{final})}$  and output with graph context

Section	Content
<b>Instruction</b>	Based on initial LLM reasoning and the subsignaling gene regulatory network identified by the subgraph generator, please identify the 100 priority genes whose knockout causes the strongest negative effect on the viability or proliferation of cell line $c_n$ in the context of disease $s'_n$ .
<b>Input</b>	<ul style="list-style-type: none"> <li>- Top 10 ranked genes with amplification from copy number data: <math>g_1'^{(n)}, \dots, g_K'^{(n)}</math></li> <li>- Top 10 ranked transcripts with high expression: <math>t_1'^{(n)}, \dots, t_K'^{(n)}</math></li> <li>- Top 10 ranked proteins from RPPA: <math>p_1'^{(n)}, \dots, p_K'^{(n)}</math></li> <li>- Disease-associated proteins from the biomedical knowledge graph: <math>\mathcal{V}_n^{(\text{sub})}</math></li> <li>- Known protein–protein/disease–protein relationships: <math>\mathcal{E}_n^{(\text{sub})}</math></li> <li>- Identified Subsignaling Gene Regulatory Network from Graph Generator</li> <li>- Involved genes in best connected subgraph: <math>\mathcal{V}_n^{(\text{conn})}</math></li> <li>- Inferred signaling cascade (edge text): <math>\mathcal{E}_n^{(\text{sub})}</math></li> </ul>
<b>Refined Reasoning</b>	<p>Based on the integrated multi-omics data and knowledge graph, I identified the 100 genes whose knockout is predicted to have the most severe negative impact on the viability or proliferation of the <math>c_n</math> cell line in <math>s'_n</math>. The prioritized gene list is as follows:</p> <ol style="list-style-type: none"> <li>1. <math>\hat{r}_{n,1}</math></li> <li>2. <math>\hat{r}_{n,2}</math></li> <li>3. <math>\hat{r}_{n,3}</math></li> <li>...</li> </ol> <p>These genes represent critical vulnerabilities for the given cell line under the disease context.</p>

distribution during the pretraining of the graph encoder  $f_G$ . We pretrained a Graph Attention Network (GAT) model  $f_G^{\text{pre}}$  with parameters  $\theta_G^{\text{pre}}$ , and selected the best-performing checkpoint based on test accuracy. The selected model achieved a training loss of 0.036, training accuracy of 99.46%, and training F1 score of 0.996. On the test set, it obtained a loss of 0.370, accuracy of 96.15%, and F1 score of 0.973, demonstrating strong generalization to both classes.

## D EXPERIMENT DETAILS

### D.1 PROMPT AND CONTEXT DESIGN FOR FINETUNING

Same as aforementioned prompt design shown in Table 8, we construct the final-stage prompt  $P_n^{(\text{final})}$ , which incorporates both the initial answering and the subgraph-based regulatory context, as detailed in Table 9. We perform five additional finetuning trials using this refined prompt format. The corresponding training are illustrated in Figure 10, where the best run (run 5) demonstrates smooth convergence and is selected as the final model checkpoint  $f_{\text{final}}$  for generating the ultimate target predictions.

### D.2 HYPERPARAMETERS

For finetuning, the model is trained for 5 epochs with a per-device batch size of 1 and a gradient accumulation step of 2, resulting in an effective batch size of 2. We use the AdamW optimizer (adamw\_torch) with a learning rate of  $1 \times 10^{-5}$ , no weight decay, and cosine learning rate scheduling with a 10% warm-up ratio. Gradient clipping is applied with a maximum gradient norm of 0.5. To support memory efficiency and large model training, we enable gradient checkpointing and use bfloat16 precision (with fp16 explicitly disabled). Training is accelerated using DeepSpeed Stage 3 parallelism, configured via an external JSON file. All experiments are conducted using 2 NVIDIA H100 GPUs, each with 80GB of memory. The DeepSpeed configuration enables ZeRO Stage 3 with both optimizer and parameter offloading to CPU, memory pinning, communication overlap, and gradient contiguity. Key parameters include automatic tuning of reduce bucket sizes and micro-batch size, sub-grouping disabled, and support for 16-bit weight gathering upon model saving. Gradient accumulation steps, clipping, learning rate, weight decay, and total training steps are all set to "auto" for adaptive scaling. Warmup is controlled via WarmupDecayLR, which

adjusts the schedule based on total steps. Model checkpoints and evaluation are performed every 67 steps, with up to 5 checkpoints retained. Logging is performed at every step. For data loading, we use 4 dataloader workers and retain all batches (i.e., `drop_last=False`). Both the initial answering finetuning and second stage finetuning for  $f_{\text{final}}$ .

To generate high-quality and biologically reasonable subgraph explanations, we implement a reinforcement-guided generator, denoted as  $\pi(\cdot)$ , which operates over sample-specific graph states and candidate sets. To enhance robustness under noisy or unstable generation dynamics, we introduce a retry-based mechanism that adaptively tunes key hyperparameters with multiple runs ( $\Psi = 6$ , and  $\psi$  denotes the current number of runs), dynamically adjusting the configuration to encourage exploration and improve convergence. Specifically, the number of training epochs is reduced from 5 to a minimum of 2, promoting faster reinitialization. Simultaneously, the learning rate is increased linearly from an initial value of  $1 \times 10^{-3}$  to  $0.001 \cdot (1 + \psi)$  to escape local minima. The maximum number of nodes per graph is reduced from 200 to 100 in steps of 25. The maximum number of graph construction steps  $i$  per rollout is similarly reduced from 50 to 20 in steps of 5. The retry mechanism enables efficient navigation of the search space while preserving biological plausibility through domain-specific priors embedded in the reward formulation. To further study the effect of sparse versus dense rewards, we vary the reward calculation frequency by introducing a step interval parameter  $s$ . Specifically, when  $s = 2$ , we compute the reward only every 2 steps during graph generation. Larger values of  $s$  produce sparser, more delayed reward signals, while smaller values (approaching  $s = 1$ ) provide denser, more immediate feedback. As shown in Table 10, GALAX demonstrates consistent stability across all reward calculation intervals. Notably, performance improves as rewards become denser (smaller  $s$ ), with optimal results achieved at  $s = 1$ , the default GALAX configuration. This trend indicates that the graph generator benefits most from immediate biological feedback provided at each step by the graph process reward model. In contrast, less frequent reward calculations (larger  $s$ ) introduce noisier, longer-horizon estimates that dilute the signal quality. Importantly, even under highly sparse reward conditions (e.g.,  $s = 10$ ), GALAX consistently outperforms the baseline L3-FT(QA) + Omics + KG, demonstrating robustness to reward sparsity.

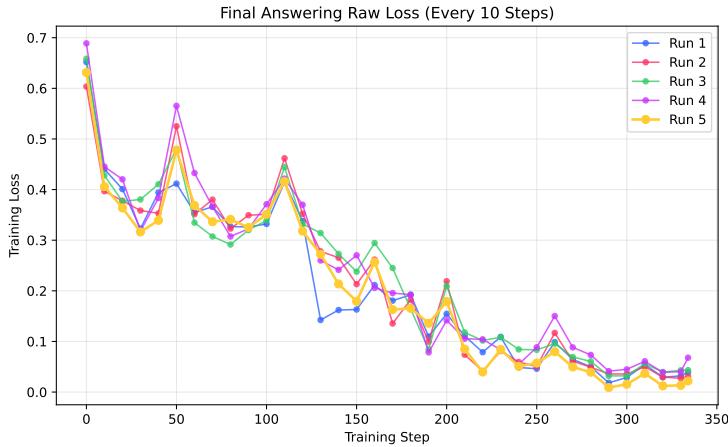
**Table 10:** Ablation study on reward density controlled by step interval  $s$

Reward Interval (s)	Precision $\uparrow$	Recall $\uparrow$	F1 Score $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
10	$0.5312 \pm 0.0044$	$0.5163 \pm 0.0019$	$0.5232 \pm 0.0031$	$0.3582 \pm 0.0035$	$0.8996 \pm 0.0050$	$0.8753 \pm 0.0044$
8	$0.5347 \pm 0.0062$	$0.5224 \pm 0.0043$	$0.5284 \pm 0.0054$	$0.3619 \pm 0.0055$	$0.9063 \pm 0.0079$	$0.8771 \pm 0.0043$
6	$0.5398 \pm 0.0051$	$0.5218 \pm 0.0036$	$0.5316 \pm 0.0044$	$0.3657 \pm 0.0044$	$0.9024 \pm 0.0079$	$0.8794 \pm 0.0040$
4	$0.5389 \pm 0.0054$	$0.5278 \pm 0.0025$	$0.5302 \pm 0.0041$	$0.3673 \pm 0.0046$	$0.9137 \pm 0.0056$	$0.8789 \pm 0.0038$
2	$0.5463 \pm 0.0057$	$0.5296 \pm 0.0032$	$0.5384 \pm 0.0040$	$0.3718 \pm 0.0043$	$0.9218 \pm 0.0048$	$0.8806 \pm 0.0033$
1 (GALAX)	<b><math>0.5472 \pm 0.0053</math></b>	<b><math>0.5332 \pm 0.0031</math></b>	<b><math>0.5399 \pm 0.0041</math></b>	<b><math>0.3726 \pm 0.0037</math></b>	<b><math>0.9249 \pm 0.0048</math></b>	<b><math>0.8815 \pm 0.0033</math></b>
L3-FT(QA) + Omics + KG	$0.5185 \pm 0.0240$	$0.4908 \pm 0.0402$	$0.5038 \pm 0.0327$	$0.3393 \pm 0.0298$	$0.8794 \pm 0.0114$	$0.8529 \pm 0.0153$

Furthermore, GALAX employs four RL components to collectively prevent reward hacking. First, the **frozen graph oncogenicity classifier (Graph)** derives reward signals from a pre-trained, frozen foundation model rather than a co-trained reward model, thereby eliminating co-evolution exploits. Second, the **schema-based rule term (Rules)** validates each action against biomedical knowledge graph ground truth to ensure that proposed edges conform to established biological constraints. Third, the **rollout-based future reward (Rollout)** evaluates each action based on its long-term consequences, preventing the generator from exploiting local reward irregularities or pursuing myopic gains. Fourth, **stepwise quality gating (Gating)** rejects actions yielding negative cumulative reward, ensuring that only biologically valid steps—those satisfying both schema rules and oncogenicity criteria—are retained. Together, these four components constitute a multi-faceted defense against reward hacking. To validate this design, we conduct an ablation study by systematically removing each component. As shown in Table 11, the robustness of GALAX does not stem from any single component but rather from the interplay of all four mechanisms. The rollout and gating modules help avoid short-sighted decisions and filter out invalid reasoning paths, while the graph supervisor and schema rules provide the underlying constraints that encourage the model to generate biologically plausible subgraphs. Notably, removing either the graph supervisor or the schema rules leads to consistent performance drops, suggesting that the frozen graph oncogenicity classifier offers reliable biological guidance, and that the schema constraints effectively block structurally incorrect actions. In sum, these mechanisms work collectively to push the generator toward producing biologically plausible subgraphs rather than outputs that score well but lack plausibility, making our RL considerably more robust than prior approaches.

**Table 11:** Ablation study of reward components

Model Variant	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
L3-FT(QA) + Omics	0.5250 $\pm$ 0.0282	0.4959 $\pm$ 0.0435	0.5094 $\pm$ 0.0365	0.3449 $\pm$ 0.0338	0.8889 $\pm$ 0.0168	0.8693 $\pm$ 0.0157
L3-FT(QA) + Omics + KG	0.5185 $\pm$ 0.0240	0.4908 $\pm$ 0.0402	0.5038 $\pm$ 0.0327	0.3393 $\pm$ 0.0298	0.8794 $\pm$ 0.0114	0.8529 $\pm$ 0.0153
GALAX w/o Graph	0.5314 $\pm$ 0.0045	0.5275 $\pm$ 0.0077	0.5336 $\pm$ 0.0053	0.3680 $\pm$ 0.0043	0.9048 $\pm$ 0.0084	0.8741 $\pm$ 0.0056
GALAX w/o Rules	0.5305 $\pm$ 0.0074	0.5154 $\pm$ 0.0115	0.5221 $\pm$ 0.0098	0.3570 $\pm$ 0.0078	0.8974 $\pm$ 0.0073	0.8746 $\pm$ 0.0042
GALAX w/o Rollout	0.5362 $\pm$ 0.0060	0.5196 $\pm$ 0.0052	0.5271 $\pm$ 0.0046	0.3628 $\pm$ 0.0043	0.9067 $\pm$ 0.0062	0.8769 $\pm$ 0.0055
GALAX w/o Gating	0.5387 $\pm$ 0.0050	0.5243 $\pm$ 0.0042	0.5308 $\pm$ 0.0024	0.3651 $\pm$ 0.0032	0.9118 $\pm$ 0.0056	0.8782 $\pm$ 0.0046
<b>GALAX (Full Model)</b>	<b>0.5472<math>\pm</math>0.0053</b>	<b>0.5332<math>\pm</math>0.0031</b>	<b>0.5399<math>\pm</math>0.0041</b>	<b>0.3726<math>\pm</math>0.0037</b>	<b>0.9249<math>\pm</math>0.0048</b>	<b>0.8815<math>\pm</math>0.0033</b>

**Figure 10:** Final language answering training loss

### D.3 BASELINE MODELS

**Multomic2Target** Multomics2Target<sup>5</sup> is a statistical and knowledge-based framework that integrates transcriptomics, proteomics, and phosphoproteomics to identify cancer-specific therapeutic targets. It leverages curated background knowledge databases and enrichment algorithms to compare tumor profiles against normal tissues, aiming to highlight targets uniquely activated in cancer. Specifically, the method filters out candidates that lack protein-level evidence or are not phosphorylated, and prioritizes those enriched in oncogenic pathways and subtype-specific signaling patterns. The workflow accounts for tumor heterogeneity by enabling both subtype-level and patient-specific analyses, thereby offering a safer and more context-aware target identification strategy. However, as a non-learning-based baseline, Multomics2Target relies purely on statistical associations, which limits its precision in complex settings. As shown in the evaluation (see Figure 11), this method demonstrates poor performance compared to other approaches, failing to accurately prioritize the most effective targets. The model takes as input the cell line’s multi-omics features and outputs a ranked report of candidate targets, but its statistical scoring alone is insufficient to capture deeper, context-specific biological relevance.

**GAT** As a graph-based baseline, we adapt a Graph Attention Network (GAT) to assess the effect of gene knockouts using multi-omics data and CRISPR-derived gene effect scores. The input graph represents a biological hierarchy, where nodes correspond to genes, transcripts, and proteins, and edges capture central dogma relationships and known molecular interactions. For each gene of interest, we simulate the perturbation process by masking the node corresponding to the gene and recursively masking its downstream elements—namely its associated transcripts and translated proteins—mimicking the transcriptional and translational disruptions observed during a real gene knockout. This masked subgraph is then passed to the GAT model based on the pretrained model  $f_G^{\text{pre}}$  (parameterized by  $\theta_G^{\text{pre}}$ ), which aggregates information from the unmasked neighborhood to predict the impact of the knockout on cellular viability, matching the gene effect values from the CRISPR dataset. This design allows us to evaluate how well the local subgraph structure and omics context can recover the observed CRISPR perturbation effect. Although this method incorporates the biological topology and omics signals through attention-weighted message passing, it lacks explicit global reasoning for disease mechanism and has low performance on CRISPR target prediction.

<sup>5</sup><https://multomics2targets.maayanlab.cloud/>

**G-Retriever** G-Retriever is a graph language framework that enables conversational interaction with large, complex graphs enriched with textual node and edge attributes. It facilitates question answering by aligning user queries in natural language with the underlying graph structure, returning both textual responses and highlighted subgraph evidence. This setup is particularly useful for scenarios where graph elements (e.g., proteins, drugs, phenotypes) are associated with rich descriptions or biomedical annotations. In our implementation, we instantiate G-Retriever with two pretrained components: (1) a language model  $f_{\text{init}}$  that has been fine-tuned on domain-specific question-answering tasks—specifically, CRISPR-target-related biomedical questions; (2) a graph encoder based on a pretrained Graph Attention Network (GAT), denoted as  $f_G^{\text{pre}}$ , parameterized by  $\theta_G^{\text{pre}}$ . The GAT is trained to encode signaling pathway graphs and capture topological and contextual dependencies between biological entities such as genes, proteins, and molecular interactions. This dual-modality initialization, textual reasoning via  $f_{\text{init}}$  and structural embedding via  $f_G^{\text{pre}}$ , allows G-Retriever to support context-aware retrieval over large graphs that exceed the LLM’s context window. Through this integration, the framework effectively aligns natural language queries with underlying graph structures, enabling biologically grounded reasoning over complex networks. While G-Retriever offers modest performance improvements over conventional GraphRAG baselines—particularly in terms of contextual relevance and retrieval accuracy—it remains limited in its ability to construct biologically interpretable subgraphs with strong explanatory power. Specifically, G-Retriever lacks a mechanism for enforcing structural coherence or pathway-level constraints, which are essential for generating subgraphs with high fidelity to known signaling or regulatory cascades. Consequently, its generated outputs often fall short of the domain-specific interpretability and performance achieved by our proposed method, which integrates graph-aware reward signals and multi-step reasoning for robust target identification.

**Other Models** RoG, SubgraphRAG, and GNN-RAG are evidence-subgraph-dependent methods that share a common assumption: a small, clean, and reliable subgraph exists from which a retriever module can be trained to provide reasoning paths to the LLM. However, this assumption breaks down in large-scale biomedical knowledge graphs, which are inherently noisier and lack validated reasoning paths. In the Target-QA setting, each question is paired with a one-hop disease–protein subgraph and a one-hop PPI neighborhood extracted from BioMedGraphica. These subgraphs are extremely noisy and large (containing thousands of nodes and edges), and lack ground-truth mechanistic subgraphs, making supervised retrieval difficult. Specifically, RoG, SubgraphRAG, and GNN-RAG require ground-truth reasoning paths for training. Since such annotations are unavailable in our setting, we employ ChatGPT 5.1 to generate candidate reasoning paths relevant to the annotated cell lines, incorporating disease context and multi-omic profiles. However, this GPT-derived supervision is inherently noisy and biologically incomplete, limiting retriever effectiveness. All baselines use the fine-tuned Llama3-8B-Instruct as the backbone LLM.

#### D.4 RESULTS

We summarize detailed per-cohort results across various cancer types from TCGA in Figure 11 and Tables 16–43, demonstrating that GALAX achieves the highest performance across nearly all evaluated metrics and cancer cohorts. M2T, serving as a baseline relying solely on multi-omics data without any structured graph or language modeling, performs poorly, underscoring the necessity of incorporating richer contextual information. Adding graph structure via a pretrained Graph Attention Network (GAT) provides some improvements, albeit limited, indicating that static graph representations alone are insufficient for capturing the complex biological interactions inherent in the data. To further elucidate the value of incorporating language models, we evaluate multiple LLaMA3 (L3) variants. Models that lack Target-specific finetuning, such as L3 combined only with omics data (L3) or with additional knowledge graphs (L3 + KG), deliver suboptimal performance, even with external knowledge enrichment. Biomedical-domain-specific finetuning (L3-FT(Med)) significantly enhances performance in identifying relevant targets, and integrating knowledge graphs yields modest incremental improvements. Nonetheless, even the variant supervised by Target-QA (L3-FT(QA) + KG) demonstrates limited additional benefit from knowledge graph integration, primarily due to its inability to dynamically generate structured subgraphs tailored to the reasoning task. The G-Retriever model, augmenting a pretrained graph attention network with retrieval mechanisms, achieves stronger performance compared to earlier baselines but remains limited by its absence of step-wise supervision and interpretability in biological reasoning. In contrast, GALAX uniquely integrates a Target-QA finetuned large language model, a pretrained Graph Attention Network, and a reinforcement learning-driven subgraph generator guided by a graph process reward

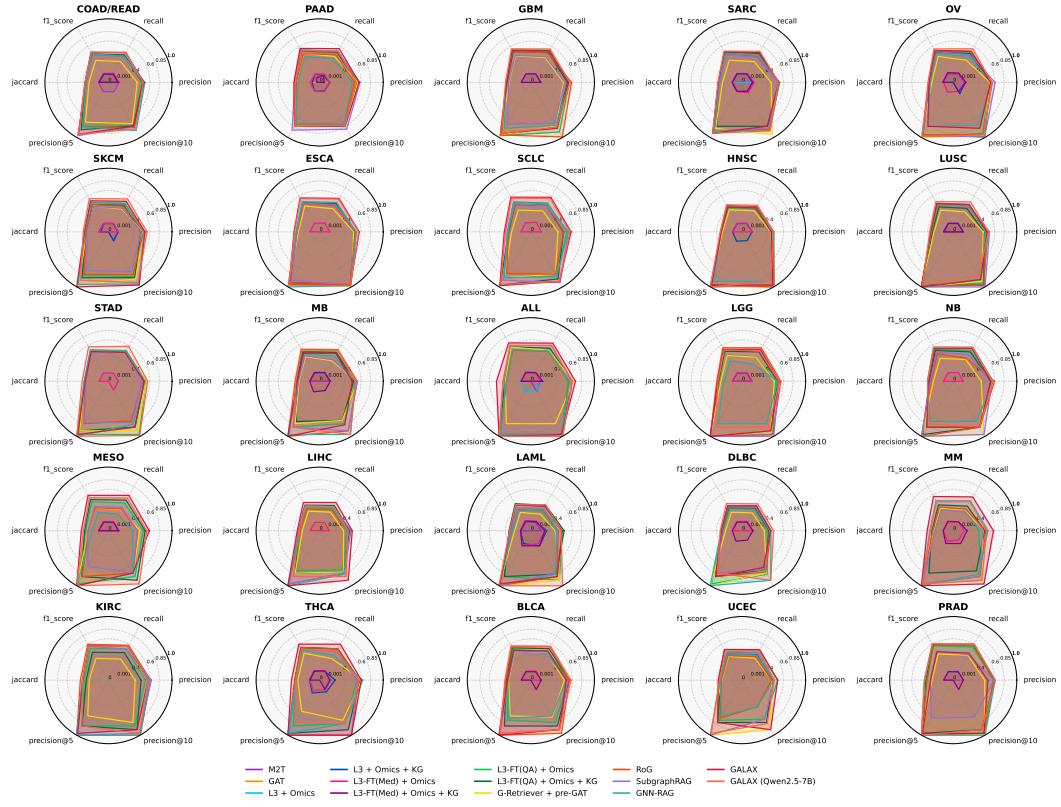


Figure 11: Model performance plots over main Target-QA datasets

model. This sophisticated architecture facilitates interpretable and biologically informed reasoning by dynamically generating structured subgraphs relevant to the biological context. Consequently, GALAX significantly surpasses all baseline models, achieving state-of-the-art performance across both overall average metrics and individual cancer cohort-specific evaluations.

### D.5 GENERALIZATION TO UNSEEN CANCER TYPES

To assess how well GALAX generalizes across cancer types, we constructed three holdout sets in which all cell lines under same TCGA codes were removed during training and used only for evaluation (see Table 12). This design forces the model to generate CRISPR target predictions for cancer types not observed during optimization, providing a direct measure of cross-cancer transfer and robustness to unseen molecular contexts.

As expected, performance decreases when entire cancer types are held out, since no fine-tuning data from those cancers is available. Nevertheless, the results in Table 13 show that GALAX maintains stable performance across the three held-out cancer groups. Precision, recall, F1, and Jaccard values remain consistent, with an average F1 of 0.4862 and Hit@10 above 0.85. While these scores are lower than the full-data GALAX model, the gap is moderate, indicating that the model retains the ability to infer relevant targets even when the TCGA code of the input cell line has never appeared during training, as long as it has been well pretrained with sufficient samples. This demonstrates that GALAX does not overfit to specific cancer labels and can transfer its reasoning across diverse cell-line backgrounds, addressing concerns regarding generalization to new cancer types.

Table 12: Held-out cancer type fold assignments

Holdout Set	Train	Test	Train/Test TCGA Codes	Hold-out Cancer TCGA Codes
<b>Holdout Set 1</b>	238 (65.6%)	125 (34.4%)	20 / 10	COAD/READ, DLBC, KIRC, LCML, LUSC PAAD, PRAD, SARC, SKCM, STAD
<b>Holdout Set 2</b>	289 (79.6%)	74 (20.4%)	21 / 9	BLCA, GBM, LAML, LGG, LIHC MESO, MM, NB, UCEC
<b>Holdout Set 3</b>	220 (60.6%)	143 (39.4%)	21 / 9	ALL, BRCA, ESCA, HNSC, LUAD MB, OV, SCLC, THCA

**Table 13:** GALAX performance on unseen cancer types under held-out evaluation

Holdout Set	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
Holdout Set 1	0.4931	0.4501	0.4699	0.3100	0.9024	0.8688
Holdout Set 2	0.5124	0.5151	0.5123	0.3465	0.8703	0.8581
Holdout Set 3	0.4832	0.4715	0.4765	0.3161	0.8993	0.8448
<b>Average</b>	0.4962	0.4789	0.4862	0.3242	0.8907	0.8572
<b>GALAX</b>	0.5472 $\pm$ 0.0053	0.5332 $\pm$ 0.0031	0.5399 $\pm$ 0.0041	0.3726 $\pm$ 0.0037	0.9249 $\pm$ 0.0048	0.8815 $\pm$ 0.0033

## D.6 EXTERNAL DATASETS

To examine generalization beyond Target-QA, we performed an external evaluation on the pediatric cancer dataset from PedDep<sup>6</sup>, which contains CRISPR-based gene effect profiles and cell-line annotations for 31 pediatric cancer samples. The data were processed following the same steps as Target-QA. Due to the limited sample size, we did not train or finetune on this dataset and instead evaluated all models in a zero-shot setting. Given that all samples are cancerous, the M2T baseline cannot be applied. As shown in Table 14, baseline graph models and language-model variants still have near-zero performance. Graph-augmented methods such as RoG, SubgraphRAG, and GNN-RAG yield moderate gains, but GALAX attains the highest precision, recall, F1, and Jaccard scores, along with notable improvements in Hit@5 and Hit@10 finetuned on Target-QA with multi-omics and disease-related proteins information. These results demonstrate that GALAX can extend its target-prediction ability to external cancer datasets without retraining, providing strong evidence of its robustness and transferability beyond the Target-QA benchmark.

**Table 14:** Model performance on PedDep cancer dataset

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
GAT	0.0005 $\pm$ 0.0008	0.0005 $\pm$ 0.0008	0.0005 $\pm$ 0.0008	0.0002 $\pm$ 0.0004	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0144 $\pm$ 0.0081	0.0114 $\pm$ 0.0049	0.0099 $\pm$ 0.0028	0.0050 $\pm$ 0.0014	0.0210 $\pm$ 0.0421	0.0109 $\pm$ 0.0218
L3-FT(Med) + Omics + KG	0.0167 $\pm$ 0.0161	0.0060 $\pm$ 0.0047	0.0068 $\pm$ 0.0049	0.0035 $\pm$ 0.0025	0.0072 $\pm$ 0.0143	0.0073 $\pm$ 0.0145
L3 + Omics	0.0054 $\pm$ 0.0108	0.0011 $\pm$ 0.0021	0.0018 $\pm$ 0.0036	0.0009 $\pm$ 0.0018	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0192 $\pm$ 0.0188	0.0032 $\pm$ 0.0022	0.0050 $\pm$ 0.0035	0.0026 $\pm$ 0.0018	0.0142 $\pm$ 0.0165	0.0214 $\pm$ 0.0247
L3-FT(QA) + Omics	0.2608 $\pm$ 0.0268	0.2605 $\pm$ 0.0252	0.2606 $\pm$ 0.0260	0.1517 $\pm$ 0.0152	0.5810 $\pm$ 0.0330	0.4524 $\pm$ 0.0502
L3-FT(QA) + Omics + KG	0.2619 $\pm$ 0.0075	0.2552 $\pm$ 0.0095	0.2584 $\pm$ 0.0086	0.1489 $\pm$ 0.0055	0.5048 $\pm$ 0.0165	0.4952 $\pm$ 0.0218
G-Retriever + pre-GAT	0.2624 $\pm$ 0.0212	0.2610 $\pm$ 0.0206	0.2617 $\pm$ 0.0209	0.1522 $\pm$ 0.0121	0.5143 $\pm$ 0.0495	0.4524 $\pm$ 0.0825
RoG	0.2730 $\pm$ 0.0050	0.2667 $\pm$ 0.0092	0.2697 $\pm$ 0.0071	0.1566 $\pm$ 0.0047	0.5143 $\pm$ 0.0286	0.4714 $\pm$ 0.0623
SubgraphRAG	0.2736 $\pm$ 0.0091	0.2690 $\pm$ 0.0095	0.2712 $\pm$ 0.0092	0.1579 $\pm$ 0.0062	0.5619 $\pm$ 0.0165	0.5286 $\pm$ 0.0571
GNN-RAG	0.2760 $\pm$ 0.0092	0.2700 $\pm$ 0.0094	0.2728 $\pm$ 0.0093	0.1589 $\pm$ 0.0064	0.5714 $\pm$ 0.0495	0.5333 $\pm$ 0.0360
<b>GALAX</b>	0.2914 $\pm$ 0.0115	0.2889 $\pm$ 0.0131	0.2901 $\pm$ 0.0123	0.1703 $\pm$ 0.0086	0.6357 $\pm$ 0.0589	0.5179 $\pm$ 0.0513
<b>GALAX (Qwen2.5-7B)</b>	0.2921 $\pm$ 0.0055	0.2895 $\pm$ 0.0050	0.2908 $\pm$ 0.0053	0.1708 $\pm$ 0.0038	0.6667 $\pm$ 0.0595	0.5381 $\pm$ 0.0705

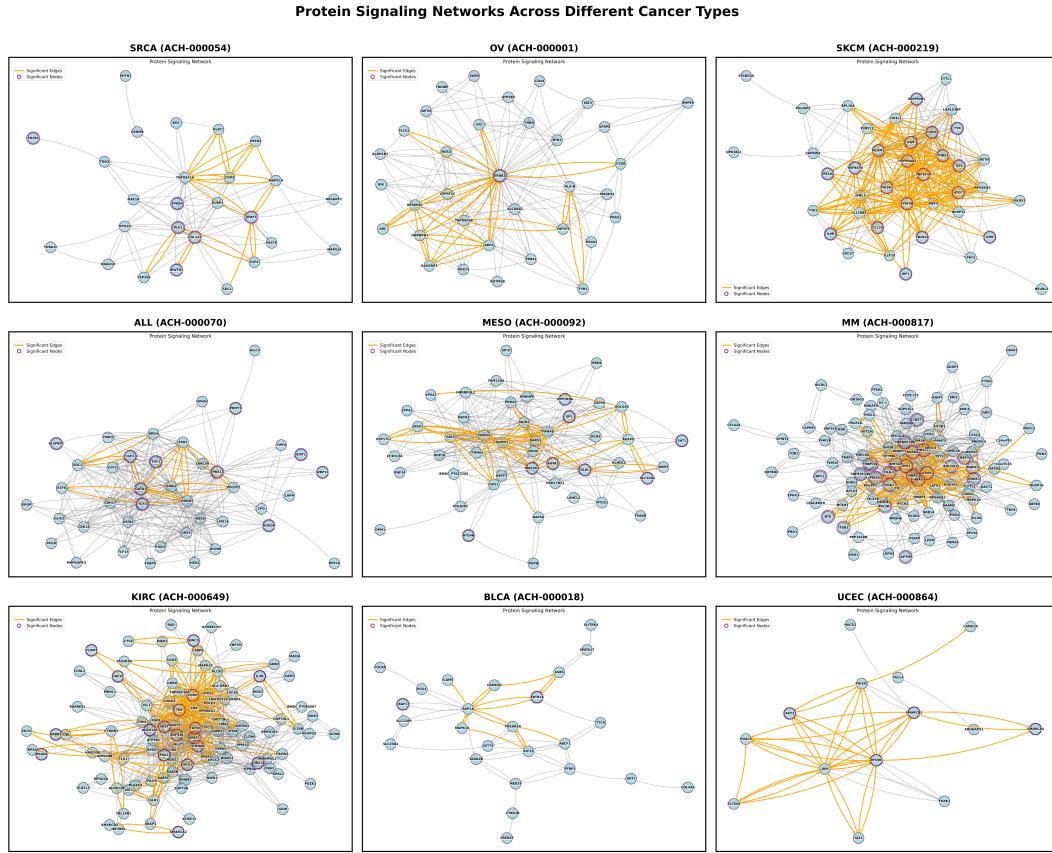
## E REINFORCEMENT LEARNING GENERATED SUBGRAPH DETAILS

Figure 12 provides a high-level overview of the explainable subgraphs of part of cell lines, highlighting key proteins and their predicted functional interactions that define the unique molecular signatures of each cell line. These interpretable network maps identify the most salient features prioritized by our model. To move beyond structural insights and assess their biological significance, we conduct enrichment analyses to evaluate whether the identified molecules are significantly over-represented in curated knowledge bases such as KEGG and WikiPathway. This analysis enables us to contextualize the salient features within relevant targets and signaling pathways. The resulting cell-line-specific interpretations bridge model-driven discovery with established biological knowledge and are detailed in the following sections.

### E.1 ENRICHMENT ANALYSIS

Enrichment analysis of **sample ACH-000054 (HT-1080)**, representing fibrosarcoma (metastasis; SARC), revealed a significant association with **Sarcoma** ( $P$ -value =  $1.625 \times 10^{-4}$ ,  $Adj.P = 4.08 \times 10^{-2}$ ) driven by **COL1A1**, **WWTR1**, **APAF1**, **PLK1**, **EPHB4**, and **SH2B1**. Pathway-level signals underscored apoptotic and stress-response programs, with strong enrichment of **Apoptosis Modulation by HSP70** (WP384;  $P$ -value =  $1.864 \times 10^{-6}$ ,  $Adj.P = 1.45 \times 10^{-4}$ ) and **Apoptosis** (WP254;  $P$ -value =  $1.73 \times 10^{-4}$ ,  $Adj.P = 5.26 \times 10^{-3}$ ), each supported by

<sup>6</sup><https://peddep.org/>



**Figure 12:** Part of explainable disease mechanism by core signaling subgraphs

**MAPK10, APAF1, and TNFRSF1A.** Additional inflammatory and mitogenic axes were indicated by **TNF alpha Signaling Pathway** (WP231;  $P$ -value =  $2.27 \times 10^{-4}$ ,  $Adj.P = 5.26 \times 10^{-3}$ ) involving **APAF1**, **PLK1**, and **TNFRSF1A**, and by the **MAPK Signaling Pathway** (WP382;  $P$ -value =  $2.70 \times 10^{-4}$ ,  $Adj.P = 5.26 \times 10^{-3}$ ) featuring **MAPK10**, **DUSP1**, **RASGRF2**, and **TNFRSF1A**. Finally, **Nanoparticle-mediated activation of receptor signaling** (WP2643;  $P$ -value =  $6.02 \times 10^{-4}$ ,  $Adj.P = 8.99 \times 10^{-3}$ ) highlighted extracellular matrix and receptor-proximal cues via **COL1A1** and **MAPK10**. Collectively, recurrent involvement of apoptosis regulators (**APAF1**, **TNFRSF1A**), MAPK components (**MAPK10**, **DUSP1**, **RASGRF2**), and proliferative drivers (**PLK1**) points to coordinated apoptotic, inflammatory, and receptor–MAPK signaling programs characteristic of fibrosarcoma biology.

For sample **ACH-000001 (NIH:OVCAR-3)**, an ovarian adenocarcinoma line (metastasis; OV), analysis indicated only a weak disease-level association with **ovarian cancer** ( $P$ -value =  $1.41 \times 10^{-1}$ ,  $Adj.P = 1.65 \times 10^{-1}$ ), primarily linked to **ERBB2**. While this signal was not significant after multiple-testing correction, pathway-level evaluation revealed several strongly dysregulated processes. The **Prolactin Signaling Pathway** (WP2037;  $P$ -value =  $2.90 \times 10^{-4}$ ,  $Adj.P = 1.49 \times 10^{-2}$ ) was prominently enriched through **RPS6KB1**, **NOS2**, and **ERBB2**, pointing to growth factor–driven oncogenic signaling. Similarly, enrichment of the **Leptin Signaling Pathway** (WP2034;  $P$ -value =  $2.90 \times 10^{-4}$ ,  $Adj.P = 1.49 \times 10^{-2}$ ) and the **ErbB Signaling Pathway** (WP673;  $P$ -value =  $4.92 \times 10^{-4}$ ,  $Adj.P = 1.69 \times 10^{-2}$ ) implicated **RPS6KB1**, **ERBB2**, and **PLCG2**, reflecting interconnected receptor tyrosine kinase and metabolic networks. Additional pathways included **IGF1–Akt signaling** (WP3850;  $P$ -value =  $1.26 \times 10^{-3}$ ,  $Adj.P = 3.09 \times 10^{-2}$ ) driven by **RPS6KB1** and **TNFSF9**, and **BDNF–TrkB Signaling** (WP3676;  $P$ -value =  $1.52 \times 10^{-3}$ ,  $Adj.P = 3.09 \times 10^{-2}$ ) via **ARC** and **RPS6KB1**. These results underscore the central role of **ERBB2**-mediated receptor activity and downstream **RPS6KB1**-linked pathways, consistent with known mechanisms of ovarian carcinoma progression.

In **sample ACH-000219 (A-375)**, representing amelanotic melanoma (primary; SKCM), enrichment testing identified a highly significant signal for **melanoma** (P-value =  $1.08 \times 10^{-6}$ , Adj.P =  $1.59 \times 10^{-4}$ ) supported by a broad gene panel including **CCL25**, **STAT5B**, **HSP90AA1**, **SERPINB1**, **VWF**, **TYR**, **THBS1**, **RELA**, **TJP1**, **GHR**, **CDH5**, **IRF1**, **TNFSF10**, **TRIM24**, **NCAM1**, **ATG7**, **TNFRSF4**, and **IL9R**. Immune and inflammatory cascades dominated the pathway enrichments, with the **NOD Pathway** (WP1433; P-value =  $1.77 \times 10^{-6}$ , Adj.P =  $2.28 \times 10^{-4}$ ) enriched via **HSP90AA1**, **NLRP12**, **MEFV**, and **RELA**, pointing to innate immune activation. Cytokine-related processes were also evident, including the **IL-4 Signaling Pathway** (WP395; P-value =  $5.42 \times 10^{-6}$ , Adj.P =  $2.65 \times 10^{-4}$ ) and the **IL-9 Signaling Pathway** (WP22; P-value =  $6.16 \times 10^{-6}$ , Adj.P =  $2.65 \times 10^{-4}$ ), with contributions from **STAT5B**, **PIK3R1**, and **IL9R**, consistent with Th2/Th9 regulation. The **Oncostatin M Signaling Pathway** (WP2374; P-value =  $1.14 \times 10^{-5}$ , Adj.P =  $3.60 \times 10^{-4}$ ) further implicated **STAT5B**, **TYK2**, **PIK3R1**, and **RELA**, linking cytokine-STAT and NF- $\kappa$ B signaling to melanoma progression. Overall, repeated enrichment of **STAT5B**, **RELA**, and **PIK3R1** across pathways highlights coordinated dysregulation of inflammatory and cytokine signaling networks in amelanotic melanoma.

Enrichment analysis of **sample ACH-000070 (697)**, representing B-cell acute lymphoblastic leukemia (ALL) with the **t(1;19)(q23;p13.3) E2A–PBX1 (TCF3–PBX1)** translocation, revealed a strong disease-level association with **acute lymphocytic leukemia** (P-value =  $7.37 \times 10^{-5}$ , Adj.P =  $8.10 \times 10^{-3}$ ) supported by genes including **IGF2**, **WBP1L**, **ELSPBP1**, **HOXD4**, **TCF3**, **THY1**, **SIRT1**, **GATA1**, and **PBX1**. Pathway enrichment further emphasized leukemia-relevant transcriptional programs: **Sudden Infant Death Syndrome (SIDS) Susceptibility Pathways** (WP706; P-value =  $3.02 \times 10^{-4}$ , Adj.P =  $1.94 \times 10^{-2}$ ) through **TCF3**, **POU2F2**, **ESR2**, and **PBX1**; and the **Wnt/β-catenin Signaling Pathway in Leukemia** (WP3658; P-value =  $1.29 \times 10^{-3}$ , Adj.P =  $4.13 \times 10^{-2}$ ) involving **AXIN2** and **TCF3**. Additional signals included the **Breast Cancer Pathway** (WP4262; P-value =  $3.85 \times 10^{-3}$ , Adj.P =  $6.06 \times 10^{-2}$ ) with **E2F3**, **AXIN2**, and **ESR2**, and the broader **Wnt Signaling Pathway** (WP363; P-value =  $5.10 \times 10^{-3}$ , Adj.P =  $6.06 \times 10^{-2}$ ) featuring **AXIN2** and **TCF3**. Collectively, these findings underscore recurrent involvement of the **TCF3–PBX1** fusion together with Wnt/β-catenin signaling and leukemogenic transcription factors, consistent with the molecular etiology of this ALL subtype.

Enrichment analysis of **sample ACH-000092 (NCI-H2452)**, representing pleural mesothelioma (metastasis; MESO), showed a modest disease-level association with **Lung Neoplasms** (P-value =  $4.965 \times 10^{-3}$ , Adj.P =  $1.89 \times 10^{-1}$ ) supported by **MAP2K1**, **SLC26A2**, **AIFM1**, **SP1**, **ATG4B**, **SAT1**, **GLUL**, and **PPP1R9B**. Pathway analysis highlighted metabolic and growth-factor-linked programs: **Amino Acid metabolism** (WP3925; P-value =  $1.20 \times 10^{-3}$ , Adj.P =  $5.63 \times 10^{-2}$ ) involving **BHMT**, **PDHA1**, and **GLUL**; the **Estrogen signaling pathway** (WP712; P-value =  $1.27 \times 10^{-3}$ , Adj.P =  $5.63 \times 10^{-2}$ ) marked by **MAP2K1** and **SP1**; and **TGF-β Signaling** (WP366; P-value =  $3.47 \times 10^{-3}$ , Adj.P =  $1.15 \times 10^{-1}$ ) featuring **MAP2K1**, **SUMO1**, and **SP1**. Together, these results suggest coordinated amino acid metabolic rewiring and transcriptional control via MAPK–SP1 and TGF-β axes in the mesothelioma context.

Enrichment analysis of **sample ACH-000817 (RPMI 8226)**, representing plasma cell myeloma (primary; MM), revealed a highly significant association with **Multiple Myeloma** (P-value =  $3.891 \times 10^{-6}$ , Adj.P =  $6.64 \times 10^{-4}$ ) supported by **ITGB1**, **KMT2D**, **CBX7**, **TNFRSF13B**, **IDH1**, **LAPTM5**, **HLA-C**, **TNFRSF10A**, **PIK3R2**, **KIR3DL1**, **BHLHA15**, **PIK3R1**, **MEFV**, **AURKA**, **NRAS**, **NUAK1**, **NPC1**, **BTK**, **KRAS**, **SGK1**, **HRAS**, **MAPK3**, and **FBXO9**. Pathway analysis further highlighted the **ErbB Signaling Pathway** (WP673; P-value =  $8.47 \times 10^{-8}$ , Adj.P =  $1.03 \times 10^{-5}$ ) involving **MAPK10**, **CAMK2D**, **NRAS**, **PIK3R2**, **KRAS**, **PIK3R1**, **HRAS**, and **MAPK3**, indicating convergence of RAS–MAPK and PI3K effector cascades downstream of ErbB family receptors—features consistent with myeloma signaling dependencies and potential therapeutic vulnerabilities.

Enrichment analysis of **sample ACH-000649 (786-O)**, representing renal cell carcinoma (primary; KIRC), revealed a highly significant association with **Conventional (Clear Cell) Renal Cell Carcinoma** (P-value =  $1.273 \times 10^{-9}$ , Adj.P =  $9.68 \times 10^{-7}$ ) supported by **TGFB1**, **IL4R**, **VCAM1**, **APAF1**, **OGG1**, **MSGN1**, **MITF**, **UNC5C**, **SMARCA2**, **TNF**, **POMC**, **FPGT**, **ORC2**, **ERBB2**, **ALDH1A1**, **PGK1**, **BIRC5**, **ZNF536**, **VHL**, and **CRYAB**. This gene set spans

hallmark features of clear cell RCC biology, including hypoxia/VHL-axis signaling (**VHL**, **PGK1**), apoptosis regulation (**APAF1**, **BIRC5**), growth factor and RTK pathways (**ERBB2**, **TGFB1**), immune–inflammatory components (**TNF**, **IL4R**, **VCAM1**), metabolic and oxidative stress responses (**ALDH1A1**, **CRYAB**), and genome maintenance/chromatin remodeling (**OGG1**, **ORC2**, **SMARCA2**), collectively underscoring a prototypical clear cell transcriptomic signature.

Enrichment analysis of **sample ACH-000018 (T24)**, representing bladder carcinoma (primary; BLCA), showed a borderline disease-level signal for **Cancer** ( $P\text{-value} = 5.755 \times 10^{-2}$ ,  $\text{Adj.P} = 1.02 \times 10^{-1}$ ) driven by **ZBTB16** and **RAF1**. Pathway-level analysis revealed significant enrichment of **Focal Adhesion** (WP306;  $P\text{-value} = 1.17 \times 10^{-4}$ ,  $\text{Adj.P} = 1.03 \times 10^{-2}$ ) involving **MAPK10**, **RAP1A**, **COL4A6**, and **RAF1**; the **ErbB Signaling Pathway** (WP673;  $P\text{-value} = 2.20 \times 10^{-4}$ ,  $\text{Adj.P} = 1.03 \times 10^{-2}$ ) featuring **MAPK10**, **CAMK2D**, and **RAF1**; and **Integrin-mediated Cell Adhesion** (WP185;  $P\text{-value} = 2.99 \times 10^{-4}$ ,  $\text{Adj.P} = 1.03 \times 10^{-2}$ ) with **MAPK10**, **RAP1A**, and **RAF1**. Collectively, recurrent involvement of **RAF1** together with **MAPK10** across adhesion and ErbB-axis pathways points to convergent RTK–MAPK and integrin signaling programs in this BLCA context.

Enrichment analysis of **sample ACH-000864 (COLO 684)**, representing endometrial adenocarcinoma (primary; UCEC), revealed a highly significant disease-level association with **Carcinoma, Small Cell** ( $P\text{-value} = 1.064 \times 10^{-6}$ ,  $\text{Adj.P} = 1.29 \times 10^{-4}$ ) supported by **SMARCB1**, **AKT3**, **MTOR**, and **SMARCA4**. Pathway analysis highlighted potent receptor tyrosine kinase and stemness programs: the **ErbB Signaling Pathway** (WP673;  $P\text{-value} = 3.38 \times 10^{-9}$ ,  $\text{Adj.P} = 4.90 \times 10^{-7}$ ) involving **JUN**, **CAMK2A**, **AKT3**, **PIK3R2**, and **MTOR**; the **EGF/EGFR Signaling Pathway** (WP437;  $P\text{-value} = 6.19 \times 10^{-8}$ ,  $\text{Adj.P} = 4.48 \times 10^{-6}$ ) through **JUN**, **GJA1**, **CAMK2A**, **PIK3R2**, and **MTOR**; and **ESC Pluripotency Pathways** (WP3931;  $P\text{-value} = 1.03 \times 10^{-6}$ ,  $\text{Adj.P} = 4.97 \times 10^{-5}$ ) marked by **JUN**, **AKT3**, **PIK3R2**, and **MTOR**. Collectively, recurrent involvement of **AKT3** and **MTOR**—together with **PIK3R2**, **JUN**, and **CAMK2A**—points to convergent EGFR/ErbB–PI3K–AKT–mTOR signaling and pluripotency-associated programs in this UCEC context.

## E.2 HUMAN EVALUATION AND LLM AS JUDGE

To further assess the biological relevance of the generated subgraphs, we conducted a comprehensive evaluation involving both human domain experts and Large Language Models (LLMs). The human evaluation panel consisted of three bioinformaticians, denoted as  $h_1$ ,  $h_2$ , and  $h_3$ . In parallel, we used two advanced LLMs as automated evaluators ( $m_1$ ,  $m_2$ ): ChatGPT-5.1 and Gemini-3.0 Pro. Both the human experts and the LLMs were provided with the generated subgraphs, corresponding gene and pathway enrichment analysis results to facilitate the assessment of biological plausibility.

**Table 15:** Human Evaluation Scores

Sample ID	TCGA Code	$h_1$	$h_2$	$h_3$	$m_1$	$m_2$	Mean $\pm$ Std
ACH-000860	LUAD	4	5	5	5	5	$4.80 \pm 0.45$
ACH-000054	SRCA	3	5	3	4	4	$3.80 \pm 0.84$
ACH-000001	OV	3	3	2	4	2	$2.80 \pm 0.84$
ACH-000219	SKCM	4	5	2	4	3	$3.60 \pm 1.14$
ACH-000070	ALL	5	5	5	4	5	$4.80 \pm 0.45$
ACH-000092	MESO	3	3	2	3	3	$2.80 \pm 0.45$
ACH-000817	MM	5	5	5	3	3	$4.20 \pm 1.10$
ACH-000649	KIRC	4	4	5	5	5	$4.60 \pm 0.55$
ACH-000018	BLCA	3	3	3	3	3	$3.00 \pm 0.00$
ACH-000864	UCEC	3	4	2	4	3	$3.20 \pm 0.84$
<b>Overall</b>		-	3.7	4.2	3.4	3.9	<b><math>3.76 \pm 0.80</math></b>

Each expert independently reviewed ten subgraph examples shown in Table 15 and assigned a score from 1 to 5 based on the degree of correspondence between the subgraph and the known biology of the associated TCGA cancer type. The scoring rubric was defined as follows:

- 5: Highly related — strong and clear match to hallmark pathways and well-established features of the cancer type.
- 4: Related — clear relationship to the cancer type but less comprehensive or slightly mixed.

- 3: Moderately related — some relevant elements present, but mixed with non-specific findings.
- 2: Not related — only weak or indirect connection to the cancer type.
- 1: Not related at all — no meaningful alignment with the biology of the cancer type.

As shown in Table 15, several samples were evaluated as related or highly related to their TCGA cancer types, indicating that the subgraphs generated by the model capture meaningful molecular features. The LUAD sample ACH-000860 received one of the highest average scores, as its subgraph highlighted pathways such as regulation of phosphatidylinositol 3-kinase signaling and positive regulation of kinase activity, involving genes like EGFR, PTK2, and EPHB4 that are central to lung adenocarcinoma biology. The ALL sample ACH-000070 was also rated highly due to the presence of TCF3 and PBX1, whose fusion is a well-established driver event in Acute Lymphoblastic Leukemia.

**Table 16:** Model Overall performance on Target-QA

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0016	0.0011	0.0013	0.0006	0.0000	0.0029
GAT	0.0006 $\pm$ 0.0000	0.0006 $\pm$ 0.0000	0.0006 $\pm$ 0.0000	0.0003 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics	0.0071 $\pm$ 0.0032	0.0013 $\pm$ 0.0002	0.0021 $\pm$ 0.0002	0.0011 $\pm$ 0.0001	0.0032 $\pm$ 0.0055	0.0021 $\pm$ 0.0037
L3 + Omics + KG	0.0125 $\pm$ 0.0032	0.0029 $\pm$ 0.0003	0.0043 $\pm$ 0.0002	0.0022 $\pm$ 0.0001	0.0085 $\pm$ 0.0037	0.0122 $\pm$ 0.0033
L3-FT(Med) + Omics	0.0179 $\pm$ 0.0045	0.0133 $\pm$ 0.0064	0.0115 $\pm$ 0.0044	0.0059 $\pm$ 0.0023	0.0116 $\pm$ 0.0097	0.0122 $\pm$ 0.0072
L3-FT(Med) + Omics + KG	0.0158 $\pm$ 0.0030	0.0058 $\pm$ 0.0011	0.0074 $\pm$ 0.0016	0.0038 $\pm$ 0.0010	0.0106 $\pm$ 0.0048	0.0132 $\pm$ 0.0040
L3-FT(QA) + Omics	0.5250 $\pm$ 0.0282	0.4959 $\pm$ 0.0435	0.5094 $\pm$ 0.0365	0.3449 $\pm$ 0.0338	0.8889 $\pm$ 0.0168	0.8693 $\pm$ 0.0157
L3-FT(QA) + Omics + KG	0.5185 $\pm$ 0.0240	0.4908 $\pm$ 0.0402	0.5038 $\pm$ 0.0327	0.3393 $\pm$ 0.0298	0.8794 $\pm$ 0.0114	0.8529 $\pm$ 0.0153
G-Retriever + pre-GAT	0.4763 $\pm$ 0.0004	0.3929 $\pm$ 0.0063	0.4286 $\pm$ 0.0044	0.2757 $\pm$ 0.0038	0.8804 $\pm$ 0.0037	0.8550 $\pm$ 0.0046
RoG	0.5248 $\pm$ 0.0134	0.4726 $\pm$ 0.0445	0.4924 $\pm$ 0.0323	0.3338 $\pm$ 0.0267	0.8593 $\pm$ 0.0318	0.8450 $\pm$ 0.0350
SubgraphRAG	0.5280 $\pm$ 0.0044	0.4617 $\pm$ 0.0027	0.4860 $\pm$ 0.0033	0.3269 $\pm$ 0.0024	0.8624 $\pm$ 0.0120	0.8476 $\pm$ 0.0167
GNN-RAG	0.5258 $\pm$ 0.0126	0.4735 $\pm$ 0.0190	0.4935 $\pm$ 0.0168	0.3345 $\pm$ 0.0134	0.8656 $\pm$ 0.0302	0.8323 $\pm$ 0.0205
GALAX	<b>0.5472<math>\pm</math>0.0053</b>	0.5332 $\pm$ 0.0031	0.5399 $\pm$ 0.0041	0.3726 $\pm$ 0.0037	<b>0.9249<math>\pm</math>0.0048</b>	0.8815 $\pm$ 0.0033
GALAX (Qwen2.5-7B)	0.5445 $\pm$ 0.0114	<b>0.5405<math>\pm</math>0.0101</b>	<b>0.5422<math>\pm</math>0.0104</b>	<b>0.3744<math>\pm</math>0.0098</b>	0.9079 $\pm$ 0.0084	<b>0.8841<math>\pm</math>0.0126</b>

**Table 17:** Model performance on LUAD

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0020	0.0014	0.0017	0.0008	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0079 $\pm$ 0.0137	0.0005 $\pm$ 0.0008	0.0009 $\pm$ 0.0016	0.0005 $\pm$ 0.0008	0.0095 $\pm$ 0.0165	0.0048 $\pm$ 0.0082
L3 + Omics + KG	0.0014 $\pm$ 0.0025	0.0010 $\pm$ 0.0016	0.0011 $\pm$ 0.0020	0.0006 $\pm$ 0.0010	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0091 $\pm$ 0.0018	0.0105 $\pm$ 0.0044	0.0079 $\pm$ 0.0022	0.0040 $\pm$ 0.0011	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0081 $\pm$ 0.0071	0.0024 $\pm$ 0.0016	0.0025 $\pm$ 0.0002	0.0013 $\pm$ 0.0001	0.0095 $\pm$ 0.0165	0.0048 $\pm$ 0.0082
L3-FT(CRISPR) + Omics	0.5201 $\pm$ 0.0408	0.4905 $\pm$ 0.0532	0.5045 $\pm$ 0.0475	0.3396 $\pm$ 0.0433	0.8476 $\pm$ 0.0165	0.8667 $\pm$ 0.0218
L3-FT(CRISPR) + Omics + KG	0.5214 $\pm$ 0.0242	0.4952 $\pm$ 0.0432	0.5073 $\pm$ 0.0343	0.3416 $\pm$ 0.0314	0.7905 $\pm$ 0.0436	0.8048 $\pm$ 0.0541
G-Retriever + pre-GAT	0.4642 $\pm$ 0.0181	0.3881 $\pm$ 0.0264	0.4204 $\pm$ 0.0233	0.2671 $\pm$ 0.0188	0.8857 $\pm$ 0.0000	0.8524 $\pm$ 0.0165
RoG	0.5213 $\pm$ 0.0227	0.4562 $\pm$ 0.0848	0.4793 $\pm$ 0.0630	0.3228 $\pm$ 0.0530	0.8095 $\pm$ 0.0436	0.8238 $\pm$ 0.0218
SubgraphRAG	0.5123 $\pm$ 0.0105	0.4448 $\pm$ 0.0386	0.4684 $\pm$ 0.0279	0.3114 $\pm$ 0.0237	0.8190 $\pm$ 0.0165	0.8238 $\pm$ 0.0082
GNN-RAG	0.5334 $\pm$ 0.0225	0.5052 $\pm$ 0.0170	0.5165 $\pm$ 0.0161	0.3563 $\pm$ 0.0127	0.7905 $\pm$ 0.0719	0.7571 $\pm$ 0.0623
GALAX	0.5345 $\pm$ 0.0185	0.5157 $\pm$ 0.0043	0.5247 $\pm$ 0.0109	0.3581 $\pm$ 0.0101	<b>0.9238<math>\pm</math>0.0436</b>	<b>0.8810<math>\pm</math>0.0082</b>
GALAX (Qwen2.5-7B)	<b>0.5475<math>\pm</math>0.0019</b>	<b>0.5462<math>\pm</math>0.0111</b>	<b>0.5465<math>\pm</math>0.0050</b>	<b>0.3778<math>\pm</math>0.0046</b>	0.9048 $\pm$ 0.0165	0.8667 $\pm$ 0.0082

**Table 18:** Model performance on BRCA

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0033 $\pm$ 0.0000	0.0033 $\pm$ 0.0000	0.0033 $\pm$ 0.0000	0.0017 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics	0.0020 $\pm$ 0.0035	0.0017 $\pm$ 0.0029	0.0018 $\pm$ 0.0032	0.0009 $\pm$ 0.0016	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0073 $\pm$ 0.0068	0.0033 $\pm$ 0.0029	0.0044 $\pm$ 0.0038	0.0022 $\pm$ 0.0019	0.0111 $\pm$ 0.0192	0.0056 $\pm$ 0.0096
L3-FT(Med) + Omics	0.0110 $\pm$ 0.0086	0.0106 $\pm$ 0.0075	0.0093 $\pm$ 0.0068	0.0047 $\pm$ 0.0035	0.0000 $\pm$ 0.0000	0.0111 $\pm$ 0.0192
L3-FT(Med) + Omics + KG	0.0149 $\pm$ 0.0057	0.0050 $\pm$ 0.0000	0.0071 $\pm$ 0.0012	0.0036 $\pm$ 0.0006	0.0000 $\pm$ 0.0000	0.0111 $\pm$ 0.0192
L3-FT(CRISPR) + Omics	0.5074 $\pm$ 0.0498	0.4856 $\pm$ 0.0570	0.4956 $\pm$ 0.0535	0.3336 $\pm$ 0.0491	<b>0.8889<math>\pm</math>0.0509</b>	0.8389 $\pm$ 0.0096
L3-FT(CRISPR) + Omics + KG	0.4856 $\pm$ 0.0395	0.4656 $\pm$ 0.0436	0.4751 $\pm$ 0.0412	0.3142 $\pm$ 0.0354	0.8778 $\pm$ 0.0192	0.8222 $\pm$ 0.0347
G-Retriever + pre-GAT	0.4414 $\pm$ 0.0099	0.3772 $\pm$ 0.0010	0.4062 $\pm$ 0.0034	0.2607 $\pm$ 0.0032	0.8667 $\pm$ 0.0000	<b>0.8667<math>\pm</math>0.0000</b>
RoG	0.4791 $\pm$ 0.0575	0.4311 $\pm$ 0.0721	0.4489 $\pm$ 0.0658	0.2999 $\pm$ 0.0528	0.7667 $\pm$ 0.0577	0.7611 $\pm$ 0.1110
SubgraphRAG	0.4708 $\pm$ 0.0317	0.3917 $\pm$ 0.0376	0.4196 $\pm$ 0.0339	0.2742 $\pm$ 0.0271	0.7556 $\pm$ 0.0839	0.7333 $\pm$ 0.1014
GNN-RAG	0.4787 $\pm$ 0.0453	0.4389 $\pm$ 0.0584	0.4543 $\pm$ 0.0531	0.3025 $\pm$ 0.0428	0.8444 $\pm$ 0.0385	0.8222 $\pm$ 0.0674
GALAX	<b>0.5608<math>\pm</math>0.0031</b>	<b>0.5533<math>\pm</math>0.0033</b>	<b>0.5569<math>\pm</math>0.0028</b>	<b>0.3886<math>\pm</math>0.0022</b>	<b>0.8889<math>\pm</math>0.0839</b>	0.8500 $\pm$ 0.0441
GALAX (Qwen2.5-7B)	0.5171 $\pm$ 0.0474	0.5206 $\pm$ 0.0419	0.5188 $\pm$ 0.0448	0.3532 $\pm$ 0.0392	0.8556 $\pm$ 0.0385	0.8000 $\pm$ 0.0764

**Table 19:** Model performance on COAD/READ

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0023	0.0017	0.0019	0.0010	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0233 $\pm$ 0.0252	0.0020 $\pm$ 0.0020	0.0035 $\pm$ 0.0033	0.0018 $\pm$ 0.0017	0.0133 $\pm$ 0.0231	0.0067 $\pm$ 0.0115
L3 + Omics + KG	0.0055 $\pm$ 0.0055	0.0033 $\pm$ 0.0031	0.0041 $\pm$ 0.0039	0.0021 $\pm$ 0.0020	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0300 $\pm$ 0.0249	0.0220 $\pm$ 0.0278	0.0199 $\pm$ 0.0214	0.0104 $\pm$ 0.0114	0.0267 $\pm$ 0.0462	0.0333 $\pm$ 0.0306
L3-FT(Med) + Omics + KG	0.0044 $\pm$ 0.0038	0.0033 $\pm$ 0.0031	0.0032 $\pm$ 0.0028	0.0016 $\pm$ 0.0014	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5011 $\pm$ 0.0335	0.4727 $\pm$ 0.0514	0.4862 $\pm$ 0.0430	0.3230 $\pm$ 0.0384	0.8400 $\pm$ 0.0800	0.8267 $\pm$ 0.0808
L3-FT(CRISPR) + Omics + KG	0.5116 $\pm$ 0.0216	0.4807 $\pm$ 0.0469	0.4950 $\pm$ 0.0351	0.3301 $\pm$ 0.0312	0.8800 $\pm$ 0.0000	0.8333 $\pm$ 0.0231
G-Retriever + pre-GAT	0.4417 $\pm$ 0.0005	0.3680 $\pm$ 0.0104	0.3995 $\pm$ 0.0073	0.2504 $\pm$ 0.0062	0.7200 $\pm$ 0.0000	0.7667 $\pm$ 0.0115
RoG	0.5184 $\pm$ 0.0462	0.4627 $\pm$ 0.0803	0.4850 $\pm$ 0.0658	0.3254 $\pm$ 0.0519	0.8533 $\pm$ 0.0611	0.8600 $\pm$ 0.0400
SubgraphRAG	0.5306 $\pm$ 0.0066	0.4627 $\pm$ 0.0854	0.4871 $\pm$ 0.0556	0.3243 $\pm$ 0.0473	<b>0.9333<math>\pm</math>0.0231</b>	0.8533 $\pm$ 0.0115
GNN-RAG	<b>0.5383<math>\pm</math>0.0306</b>	0.4627 $\pm$ 0.0463	0.4909 $\pm$ 0.0404	0.3295 $\pm$ 0.0352	0.8533 $\pm$ 0.0462	<b>0.8867<math>\pm</math>0.0115</b>
GALAX	0.5272 $\pm$ 0.0143	0.5120 $\pm$ 0.0087	<b>0.5192<math>\pm</math>0.0103</b>	<b>0.3512<math>\pm</math>0.0098</b>	0.9200 $\pm$ 0.0693	0.8400 $\pm$ 0.0200
GALAX (Qwen2.5-7B)	0.5177 $\pm$ 0.0124	<b>0.5187<math>\pm</math>0.0061</b>	0.5177 $\pm$ 0.0083	0.3498 $\pm$ 0.0076	0.9200 $\pm$ 0.0800	0.8733 $\pm$ 0.0503

**Table 20:** Model performance on PAAD

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0017	0.0013	0.0014	0.0008	0.0000	0.0000
GAT	0.0025 $\pm$ 0.0000	0.0025 $\pm$ 0.0000	0.0025 $\pm$ 0.0000	0.0013 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0139 $\pm$ 0.0241	0.0008 $\pm$ 0.0014	0.0016 $\pm$ 0.0027	0.0008 $\pm$ 0.0014	0.0167 $\pm$ 0.0289	0.0083 $\pm$ 0.0144
L3-FT(Med) + Omics	0.0051 $\pm$ 0.0061	0.0050 $\pm$ 0.0043	0.0039 $\pm$ 0.0036	0.0020 $\pm$ 0.0018	0.0167 $\pm$ 0.0289	0.0083 $\pm$ 0.0144
L3-FT(Med) + Omics + KG	0.0005 $\pm$ 0.0008	0.0008 $\pm$ 0.0014	0.0006 $\pm$ 0.0010	0.0003 $\pm$ 0.0005	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5649 $\pm$ 0.0559	0.5358 $\pm$ 0.0686	0.5492 $\pm$ 0.0605	0.3811 $\pm$ 0.0584	0.8000 $\pm$ 0.0500	0.8417 $\pm$ 0.0577
L3-FT(CRISPR) + Omics + KG	0.5499 $\pm$ 0.0570	0.5183 $\pm$ 0.0794	0.5330 $\pm$ 0.0683	0.3657 $\pm$ 0.0650	0.8333 $\pm$ 0.0764	0.8333 $\pm$ 0.0520
G-Retriever + pre-GAT	0.5594 $\pm$ 0.0114	0.4567 $\pm$ 0.0058	0.5016 $\pm$ 0.0009	0.3387 $\pm$ 0.0007	0.8500 $\pm$ 0.0000	0.8500 $\pm$ 0.0000
RoG	0.5196 $\pm$ 0.1322	0.4842 $\pm$ 0.1461	0.4973 $\pm$ 0.1406	0.3462 $\pm$ 0.1107	0.7667 $\pm$ 0.1528	0.7833 $\pm$ 0.1181
SubgraphRAG	<b>0.5873<math>\pm</math>0.0050</b>	0.5308 $\pm$ 0.0029	0.5518 $\pm$ 0.0037	0.3869 $\pm$ 0.0028	<b>0.8833<math>\pm</math>0.0289</b>	<b>0.8750<math>\pm</math>0.0000</b>
GNN-RAG	0.4923 $\pm$ 0.0762	0.4283 $\pm$ 0.0204	0.4522 $\pm$ 0.0411	0.3068 $\pm$ 0.0248	0.7667 $\pm$ 0.1443	0.7917 $\pm$ 0.1233
GALAX	0.5771 $\pm$ 0.0174	<b>0.5708<math>\pm</math>0.0118</b>	<b>0.5739<math>\pm</math>0.0145</b>	<b>0.4044<math>\pm</math>0.0154</b>	0.8333 $\pm$ 0.0289	0.8500 $\pm$ 0.0000
GALAX (Qwen2.5-7B)	0.5394 $\pm$ 0.0417	0.5383 $\pm$ 0.0427	0.5387 $\pm$ 0.0422	0.3718 $\pm$ 0.0402	0.8167 $\pm$ 0.0289	0.8333 $\pm$ 0.0144

**Table 21:** Model performance on GBM

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0056 $\pm$ 0.0096	0.0011 $\pm$ 0.0019	0.0019 $\pm$ 0.0032	0.0009 $\pm$ 0.0016	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0069 $\pm$ 0.0045	0.0133 $\pm$ 0.0088	0.0090 $\pm$ 0.0059	0.0046 $\pm$ 0.0030	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0107 $\pm$ 0.0158	0.0056 $\pm$ 0.0038	0.0064 $\pm$ 0.0074	0.0032 $\pm$ 0.0037	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5641 $\pm$ 0.0310	0.5333 $\pm$ 0.0462	0.5480 $\pm$ 0.0387	0.3794 $\pm$ 0.0364	0.9333 $\pm$ 0.0667	0.9000 $\pm$ 0.0333
L3-FT(CRISPR) + Omics + KG	0.5391 $\pm$ 0.0213	0.5178 $\pm$ 0.0164	0.5281 $\pm$ 0.0181	0.3596 $\pm$ 0.0165	<b>0.9556<math>\pm</math>0.0385</b>	0.8444 $\pm$ 0.0694
G-Retriever + pre-GAT	0.4906 $\pm$ 0.0075	0.4411 $\pm$ 0.0038	0.4642 $\pm$ 0.0010	0.3037 $\pm$ 0.0000	<b>0.9556<math>\pm</math>0.0385</b>	0.8333 $\pm$ 0.0577
RoG	<b>0.5910<math>\pm</math>0.0281</b>	0.5078 $\pm$ 0.1165	0.5380 $\pm$ 0.0853	0.3734 $\pm$ 0.0756	<b>0.9556<math>\pm</math>0.0770</b>	<b>0.9444<math>\pm</math>0.0509</b>
SubgraphRAG	0.5066 $\pm$ 0.0597	0.4544 $\pm$ 0.1001	0.4707 $\pm$ 0.0884	0.3176 $\pm$ 0.0729	0.7778 $\pm$ 0.1925	0.7556 $\pm$ 0.2117
GNN-RAG	0.5567 $\pm$ 0.0222	0.5478 $\pm$ 0.0212	0.5521 $\pm$ 0.0219	0.3855 $\pm$ 0.0141	0.8667 $\pm$ 0.2309	0.8111 $\pm$ 0.1540
GALAX	0.5527 $\pm$ 0.0021	0.5422 $\pm$ 0.0139	0.5473 $\pm$ 0.0081	0.3769 $\pm$ 0.0079	0.9111 $\pm$ 0.0770	0.8667 $\pm$ 0.0000
GALAX (Qwen2.5-7B)	0.5715 $\pm$ 0.0228	<b>0.5656<math>\pm</math>0.0269</b>	<b>0.5685<math>\pm</math>0.0249</b>	<b>0.3985<math>\pm</math>0.0236</b>	0.9111 $\pm$ 0.0770	0.8556 $\pm$ 0.0385

**Table 22:** Model performance on SARC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0047	0.0033	0.0039	0.0019	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0111 $\pm$ 0.0192	0.0033 $\pm$ 0.0058	0.0051 $\pm$ 0.0089	0.0026 $\pm$ 0.0045	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0271 $\pm$ 0.0062	0.0111 $\pm$ 0.0019	0.0157 $\pm$ 0.0029	0.0080 $\pm$ 0.0015	0.0222 $\pm$ 0.0385	0.0222 $\pm$ 0.0192
L3-FT(Med) + Omics	0.0689 $\pm$ 0.0631	0.0167 $\pm$ 0.0115	0.0229 $\pm$ 0.0180	0.0119 $\pm$ 0.0094	0.0222 $\pm$ 0.0385	0.0667 $\pm$ 0.0577
L3-FT(Med) + Omics + KG	0.0270 $\pm$ 0.0108	0.0089 $\pm$ 0.0019	0.0120 $\pm$ 0.0009	0.0061 $\pm$ 0.0005	0.0222 $\pm$ 0.0385	0.0222 $\pm$ 0.0192
L3-FT(CRISPR) + Omics	0.5347 $\pm$ 0.0167	0.5089 $\pm$ 0.0353	0.5211 $\pm$ 0.0266	0.3547 $\pm$ 0.0253	<b>0.9111<math>\pm</math>0.0385</b>	0.8778 $\pm$ 0.0192
L3-FT(CRISPR) + Omics + KG	0.5423 $\pm$ 0.0207	0.5022 $\pm$ 0.0539	0.5207 $\pm$ 0.0383	0.3570 $\pm$ 0.0344	0.8444 $\pm$ 0.0385	0.8444 $\pm$ 0.0385
G-Retriever + pre-GAT	0.4461 $\pm$ 0.0308	0.3667 $\pm$ 0.0115	0.4009 $\pm$ 0.0055	0.2518 $\pm$ 0.0051	0.8667 $\pm$ 0.0000	<b>0.9222<math>\pm</math>0.0192</b>
RoG	<b>0.5533<math>\pm</math>0.0201</b>	0.4589 $\pm$ 0.0740	0.4922 $\pm$ 0.0433	0.3324 $\pm$ 0.0375	0.8889 $\pm$ 0.0385	0.8778 $\pm$ 0.0192
SubgraphRAG	0.5365 $\pm$ 0.0277	0.4456 $\pm$ 0.0168	0.4772 $\pm$ 0.0188	0.3175 $\pm$ 0.0160	0.8889 $\pm$ 0.1018	0.8667 $\pm$ 0.0882
GNN-RAG	0.5485 $\pm$ 0.0097	0.4778 $\pm$ 0.0587	0.5036 $\pm$ 0.0423	0.3426 $\pm$ 0.0403	0.8889 $\pm$ 0.0385	0.8889 $\pm$ 0.0385
GALAX	0.5414 $\pm$ 0.0143	0.5144 $\pm$ 0.0168	0.5271 $\pm$ 0.0157	0.3602 $\pm$ 0.0165	<b>0.9111<math>\pm</math>0.1018</b>	0.8444 $\pm$ 0.0509
GALAX (Qwen2.5-7B)	0.5383 $\pm$ 0.0096	<b>0.5311<math>\pm</math>0.0084</b>	<b>0.5347<math>\pm</math>0.0090</b>	<b>0.3674<math>\pm</math>0.0088</b>	<b>0.9111<math>\pm</math>0.1018</b>	0.8889 $\pm$ 0.0694

**Table 23:** Model performance on OV

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0241 $\pm$ 0.0251	0.0044 $\pm$ 0.0051	0.0075 $\pm$ 0.0085	0.0038 $\pm$ 0.0043	0.0000 $\pm$ 0.0000	0.0111 $\pm$ 0.0192
L3 + Omics + KG	0.0660 $\pm$ 0.0667	0.0111 $\pm$ 0.0102	0.0183 $\pm$ 0.0161	0.0093 $\pm$ 0.0082	0.0000 $\pm$ 0.0000	0.0889 $\pm$ 0.0770
L3-FT(Med) + Omics	0.0356 $\pm$ 0.0412	0.0100 $\pm$ 0.0067	0.0102 $\pm$ 0.0072	0.0051 $\pm$ 0.0036	0.0444 $\pm$ 0.0770	0.0222 $\pm$ 0.0385
L3-FT(Med) + Omics + KG	0.0797 $\pm$ 0.0496	0.0344 $\pm$ 0.0280	0.0457 $\pm$ 0.0375	0.0253 $\pm$ 0.0226	0.0000 $\pm$ 0.0000	0.0333 $\pm$ 0.0333
L3-FT(CRISPR) + Omics	0.5469 $\pm$ 0.0274	0.5167 $\pm$ 0.0406	0.5307 $\pm$ 0.0325	0.3631 $\pm$ 0.0295	<b>1.0000<math>\pm</math>0.0000</b>	0.9111 $\pm$ 0.0694
L3-FT(CRISPR) + Omics + KG	0.5439 $\pm$ 0.0066	0.4989 $\pm$ 0.0195	0.5200 $\pm$ 0.0123	0.3542 $\pm$ 0.0095	0.9556 $\pm$ 0.0385	0.9444 $\pm$ 0.0385
G-Retriever + pre-GAT	0.5413 $\pm$ 0.0525	0.4033 $\pm$ 0.0173	0.4575 $\pm$ 0.0067	0.2991 $\pm$ 0.0041	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9667<math>\pm</math>0.0000</b>
RoG	0.5708 $\pm$ 0.0158	0.5289 $\pm$ 0.0523	0.5453 $\pm$ 0.0343	0.3774 $\pm$ 0.0317	0.9556 $\pm$ 0.0385	0.9222 $\pm$ 0.0385
SubgraphRAG	<b>0.6143<math>\pm</math>0.0414</b>	0.4833 $\pm$ 0.0384	0.5289 $\pm$ 0.0215	0.3650 $\pm$ 0.0185	0.9333 $\pm$ 0.0667	0.9444 $\pm$ 0.0385
GNN-RAG	0.5687 $\pm$ 0.0300	0.4600 $\pm$ 0.0145	0.5026 $\pm$ 0.0061	0.3380 $\pm$ 0.0036	0.8444 $\pm$ 0.1540	0.8111 $\pm$ 0.1540
GALAX	0.5489 $\pm$ 0.0176	0.5344 $\pm$ 0.0259	0.5413 $\pm$ 0.0219	0.3714 $\pm$ 0.0206	0.8444 $\pm$ 0.0385	0.8667 $\pm$ 0.0333
GALAX (Qwen2.5-7B)	0.5643 $\pm$ 0.0190	<b>0.5689<math>\pm</math>0.0550</b>	<b>0.5654<math>\pm</math>0.0355</b>	<b>0.3957<math>\pm</math>0.0349</b>	0.9778 $\pm$ 0.0385	0.9556 $\pm$ 0.0192

**Table 24:** Model performance on SKCM

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0055	0.0033	0.0041	0.0021	0.0000	0.0333
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0028 $\pm$ 0.0049	0.0011 $\pm$ 0.0019	0.0016 $\pm$ 0.0028	0.0008 $\pm$ 0.0014	0.0000 $\pm$ 0.0000	0.0111 $\pm$ 0.0192
L3-FT(Med) + Omics	0.0047 $\pm$ 0.0041	0.0089 $\pm$ 0.0069	0.0061 $\pm$ 0.0052	0.0031 $\pm$ 0.0026	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.5145 $\pm$ 0.0437	0.4922 $\pm$ 0.0455	0.5028 $\pm$ 0.0443	0.3399 $\pm$ 0.0401	0.8444 $\pm$ 0.1388	0.8556 $\pm$ 0.0385
L3-FT(CRISPR) + Omics + KG	0.4871 $\pm$ 0.0529	0.4678 $\pm$ 0.0626	0.4771 $\pm$ 0.0579	0.3164 $\pm$ 0.0514	0.8667 $\pm$ 0.0667	0.8667 $\pm$ 0.0667
G-Retriever + pre-GAT	0.5123 $\pm$ 0.0545	0.4100 $\pm$ 0.0058	0.4543 $\pm$ 0.0258	0.2958 $\pm$ 0.0222	0.9333 $\pm$ 0.0000	0.8778 $\pm$ 0.0192
RoG	0.5201 $\pm$ 0.0255	0.4522 $\pm$ 0.0681	0.4797 $\pm$ 0.0520	0.3179 $\pm$ 0.0461	0.8000 $\pm$ 0.2309	0.8111 $\pm$ 0.1895
SubgraphRAG	0.4929 $\pm$ 0.0245	0.4222 $\pm$ 0.0403	0.4502 $\pm$ 0.0304	0.2929 $\pm$ 0.0260	0.6889 $\pm$ 0.2143	0.7222 $\pm$ 0.2143
GNN-RAG	0.5298 $\pm$ 0.0203	0.4956 $\pm$ 0.0367	0.5092 $\pm$ 0.0257	0.3447 $\pm$ 0.0229	0.9556 $\pm$ 0.0385	0.9222 $\pm$ 0.0385
GALAX	0.5385 $\pm$ 0.0116	0.5233 $\pm$ 0.0200	0.5307 $\pm$ 0.0160	0.3638 $\pm$ 0.0152	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9667<math>\pm</math>0.0000</b>
GALAX (Qwen2.5-7B)	<b>0.5651<math>\pm</math>0.0206</b>	<b>0.5589<math>\pm</math>0.0280</b>	<b>0.5619<math>\pm</math>0.0245</b>	<b>0.3938<math>\pm</math>0.0213</b>	0.8889 $\pm$ 0.0770	0.9111 $\pm$ 0.0962

**Table 25:** Model performance on ESCA

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0143 $\pm$ 0.0036	0.0178 $\pm$ 0.0038	0.0139 $\pm$ 0.0021	0.0071 $\pm$ 0.0011	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.5210 $\pm$ 0.0590	0.4989 $\pm$ 0.0691	0.5093 $\pm$ 0.0637	0.3438 $\pm$ 0.0577	0.9333 $\pm$ 0.0667	0.9333 $\pm$ 0.0000
L3-FT(CRISPR) + Omics + KG	0.5146 $\pm$ 0.0187	0.4878 $\pm$ 0.0051	0.5003 $\pm$ 0.0074	0.3346 $\pm$ 0.0059	0.9111 $\pm$ 0.0385	0.9333 $\pm$ 0.0000
G-Retriever + pre-GAT	0.5083 $\pm$ 0.0142	0.4156 $\pm$ 0.0038	0.4556 $\pm$ 0.0076	0.2959 $\pm$ 0.0064	0.9778 $\pm$ 0.0385	0.9444 $\pm$ 0.0385
RoG	0.5275 $\pm$ 0.0188	0.4622 $\pm$ 0.0534	0.4877 $\pm$ 0.0400	0.3262 $\pm$ 0.0358	0.9111 $\pm$ 0.0770	0.9556 $\pm$ 0.0385
SubgraphRAG	0.5373 $\pm$ 0.0043	0.4744 $\pm$ 0.0372	0.4990 $\pm$ 0.0259	0.3385 $\pm$ 0.0210	0.8889 $\pm$ 0.0385	0.9333 $\pm$ 0.0000
GNN-RAG	0.5405 $\pm$ 0.0036	0.5067 $\pm$ 0.0404	0.5198 $\pm$ 0.0283	0.3548 $\pm$ 0.0243	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9778<math>\pm</math>0.0192</b>
GALAX	<b>0.5775<math>\pm</math>0.0204</b>	<b>0.5667<math>\pm</math>0.0260</b>	<b>0.5720<math>\pm</math>0.0233</b>	<b>0.4031<math>\pm</math>0.0246</b>	0.9556 $\pm$ 0.0385	0.9667 $\pm$ 0.0333
GALAX (Qwen2.5-7B)	0.5706 $\pm$ 0.0228	0.5567 $\pm$ 0.0145	0.5634 $\pm$ 0.0173	0.3936 $\pm$ 0.0156	0.9333 $\pm$ 0.1155	0.9333 $\pm$ 0.0882

**Table 26:** Model performance on SCLC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0184 $\pm$ 0.0212	0.0389 $\pm$ 0.0448	0.0249 $\pm$ 0.0286	0.0131 $\pm$ 0.0153	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.5337 $\pm$ 0.0349	0.4989 $\pm$ 0.0479	0.5146 $\pm$ 0.0378	0.3476 $\pm$ 0.0341	0.8667 $\pm$ 0.1155	0.8556 $\pm$ 0.0694
L3-FT(CRISPR) + Omics + KG	0.4900 $\pm$ 0.0597	0.4667 $\pm$ 0.0677	0.4778 $\pm$ 0.0638	0.3160 $\pm$ 0.0555	0.9333 $\pm$ 0.0667	0.8778 $\pm$ 0.0694
G-Retriever + pre-GAT	0.4203 $\pm$ 0.0263	0.3667 $\pm$ 0.0115	0.3910 $\pm$ 0.0180	0.2434 $\pm$ 0.0140	0.8667 $\pm$ 0.0000	0.8000 $\pm$ 0.0000
RoG	0.4971 $\pm$ 0.1497	0.4689 $\pm$ 0.1375	0.4786 $\pm$ 0.1399	0.3313 $\pm$ 0.1074	0.7778 $\pm$ 0.2694	0.8000 $\pm$ 0.2028
SubgraphRAG	0.5516 $\pm$ 0.0586	0.4578 $\pm$ 0.1711	0.4917 $\pm$ 0.1315	0.3346 $\pm$ 0.1131	0.8889 $\pm$ 0.1388	<b>0.9111<math>\pm</math>0.1018</b>
GNN-RAG	0.5503 $\pm$ 0.0732	0.4844 $\pm$ 0.1771	0.5096 $\pm$ 0.1386	0.3512 $\pm$ 0.1203	0.8667 $\pm$ 0.0000	0.8556 $\pm$ 0.0192
GALAX	<b>0.5850<math>\pm</math>0.0110</b>	<b>0.5822<math>\pm</math>0.0102</b>	<b>0.5836<math>\pm</math>0.0105</b>	<b>0.4143<math>\pm</math>0.0102</b>	<b>0.9778<math>\pm</math>0.0385</b>	0.9000 $\pm$ 0.0333
GALAX (Qwen2.5-7B)	0.5706 $\pm$ 0.0190	0.5667 $\pm$ 0.0176	0.5686 $\pm$ 0.0182	0.3988 $\pm$ 0.0189	0.9556 $\pm$ 0.0385	0.9000 $\pm$ 0.0333

**Table 27:** Model performance on HNSC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0128 $\pm$ 0.0222	0.0017 $\pm$ 0.0029	0.0029 $\pm$ 0.0051	0.0015 $\pm$ 0.0026	0.0333 $\pm$ 0.0577	0.0167 $\pm$ 0.0289
L3-FT(Med) + Omics	0.0082 $\pm$ 0.0036	0.0100 $\pm$ 0.0050	0.0079 $\pm$ 0.0016	0.0040 $\pm$ 0.0008	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.4374 $\pm$ 0.0214	0.4167 $\pm$ 0.0208	0.4267 $\pm$ 0.0208	0.2718 $\pm$ 0.0174	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9667<math>\pm</math>0.0577</b>
L3-FT(CRISPR) + Omics + KG	0.4787 $\pm$ 0.0445	0.4450 $\pm$ 0.0676	0.4606 $\pm$ 0.0565	0.3006 $\pm$ 0.0470	0.9667 $\pm$ 0.0577	0.9500 $\pm$ 0.0000
G-Retriever + pre-GAT	0.4413 $\pm$ 0.0030	0.3700 $\pm$ 0.0000	0.4021 $\pm$ 0.0009	0.2518 $\pm$ 0.0008	<b>1.0000<math>\pm</math>0.0000</b>	0.9500 $\pm$ 0.0000
RoG	0.4718 $\pm$ 0.0254	0.4400 $\pm$ 0.0673	0.4529 $\pm$ 0.0509	0.2952 $\pm$ 0.0404	0.9667 $\pm$ 0.0577	0.9500 $\pm$ 0.0500
SubgraphRAG	0.4648 $\pm$ 0.0120	0.4233 $\pm$ 0.0635	0.4398 $\pm$ 0.0450	0.2833 $\pm$ 0.0355	<b>1.0000<math>\pm</math>0.0000</b>	0.9333 $\pm$ 0.0289
GNN-RAG	0.4655 $\pm$ 0.0351	0.4317 $\pm$ 0.0813	0.4455 $\pm$ 0.0631	0.2890 $\pm$ 0.0506	0.9000 $\pm$ 0.0000	0.9000 $\pm$ 0.0000
GALAX	0.4482 $\pm$ 0.0229	0.4283 $\pm$ 0.0208	0.4380 $\pm$ 0.0218	0.2822 $\pm$ 0.0177	0.9333 $\pm$ 0.0577	<b>0.9667<math>\pm</math>0.0289</b>
GALAX (Qwen2.5-7B)	<b>0.4787<math>\pm</math>0.0082</b>	<b>0.4667<math>\pm</math>0.0153</b>	<b>0.4722<math>\pm</math>0.0070</b>	<b>0.3095<math>\pm</math>0.0064</b>	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0866

**Table 28:** Model performance on LUSC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0222 $\pm$ 0.0192	0.0067 $\pm$ 0.0058	0.0103 $\pm$ 0.0089	0.0052 $\pm$ 0.0045	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0062 $\pm$ 0.0068	0.0050 $\pm$ 0.0050	0.0055 $\pm$ 0.0058	0.0028 $\pm$ 0.0029	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0165 $\pm$ 0.0106	0.0083 $\pm$ 0.0058	0.0085 $\pm$ 0.0010	0.0043 $\pm$ 0.0005	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0094 $\pm$ 0.0082	0.0033 $\pm$ 0.0029	0.0049 $\pm$ 0.0043	0.0025 $\pm$ 0.0021	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5208 $\pm$ 0.0525	0.4767 $\pm$ 0.0810	0.4961 $\pm$ 0.0677	0.3320 $\pm$ 0.0618	0.9667 $\pm$ 0.0577	0.9167 $\pm$ 0.0577
L3-FT(CRISPR) + Omics + KG	0.4656 $\pm$ 0.0253	0.4217 $\pm$ 0.0584	0.4420 $\pm$ 0.0427	0.2844 $\pm$ 0.0355	0.9333 $\pm$ 0.1155	0.9333 $\pm$ 0.0577
G-Retriever + pre-GAT	0.4563 $\pm$ 0.0579	0.3417 $\pm$ 0.0318	0.3902 $\pm$ 0.0418	0.2433 $\pm$ 0.0313	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0000
RoG	0.5038 $\pm$ 0.0245	0.4550 $\pm$ 0.0954	0.4717 $\pm$ 0.0730	0.3123 $\pm$ 0.0583	<b>1.0000<math>\pm</math>0.0000</b>	0.9667 $\pm$ 0.0289
SubgraphRAG	<b>0.5360<math>\pm</math>0.0133</b>	0.4550 $\pm$ 0.0433	0.4869 $\pm$ 0.0223	0.3237 $\pm$ 0.0186	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GNN-RAG	0.4854 $\pm$ 0.0257	0.4717 $\pm$ 0.0375	0.4782 $\pm$ 0.0319	0.3164 $\pm$ 0.0263	<b>1.0000<math>\pm</math>0.0000</b>	0.8833 $\pm$ 0.0577
GALAX	0.5046 $\pm$ 0.0833	0.4783 $\pm$ 0.0808	0.4909 $\pm$ 0.0819	0.3285 $\pm$ 0.0726	<b>1.0000<math>\pm</math>0.0000</b>	0.8833 $\pm$ 0.0289
GALAX (Qwen2.5-7B)	0.5237 $\pm$ 0.0212	<b>0.5150<math>\pm</math>0.0180</b>	<b>0.5193<math>\pm</math>0.0196</b>	<b>0.3509<math>\pm</math>0.0178</b>	0.9333 $\pm$ 0.0577	0.8500 $\pm$ 0.0500

**Table 29:** Model performance on STAD

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0057 $\pm$ 0.0100	0.0017 $\pm$ 0.0029	0.0026 $\pm$ 0.0045	0.0013 $\pm$ 0.0023	0.0000 $\pm$ 0.0000	0.0167 $\pm$ 0.0289
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5528 $\pm$ 0.0311	0.5200 $\pm$ 0.0409	0.5353 $\pm$ 0.0325	0.3684 $\pm$ 0.0309	0.9333 $\pm$ 0.0577	0.9333 $\pm$ 0.0289
L3-FT(CRISPR) + Omics + KG	0.5305 $\pm$ 0.0231	0.5000 $\pm$ 0.0427	0.5144 $\pm$ 0.0332	0.3486 $\pm$ 0.0322	0.9000 $\pm$ 0.0000	0.8667 $\pm$ 0.0289
G-Retriever + pre-GAT	0.5574 $\pm$ 0.0091	0.4883 $\pm$ 0.0029	0.5193 $\pm$ 0.0024	0.3550 $\pm$ 0.0043	0.9000 $\pm$ 0.0000	0.9000 $\pm$ 0.0000
RoG	0.5099 $\pm$ 0.0498	0.4950 $\pm$ 0.0606	0.5020 $\pm$ 0.0556	0.3392 $\pm$ 0.0475	0.8000 $\pm$ 0.2646	0.7167 $\pm$ 0.2021
SubgraphRAG	0.5044 $\pm$ 0.0443	0.4950 $\pm$ 0.0606	0.4993 $\pm$ 0.0531	0.3353 $\pm$ 0.0437	0.7667 $\pm$ 0.2309	0.7667 $\pm$ 0.2309
GNN-RAG	0.5387 $\pm$ 0.0068	0.5267 $\pm$ 0.0029	0.5325 $\pm$ 0.0047	0.3680 $\pm$ 0.0037	<b>1.0000<math>\pm</math>0.0000</b>	0.8167 $\pm$ 0.0577
GALAX	0.5298 $\pm$ 0.0331	0.4983 $\pm$ 0.0404	0.5134 $\pm$ 0.0371	0.3465 $\pm$ 0.0338	0.9667 $\pm$ 0.0577	0.8667 $\pm$ 0.0289
GALAX (Qwen2.5-7B)	0.5731 $\pm$ 0.0132	<b>0.5883<math>\pm</math>0.0029</b>	<b>0.5803<math>\pm</math>0.0055</b>	<b>0.4117<math>\pm</math>0.0057</b>	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9833<math>\pm</math>0.0289</b>

**Table 30:** Model performance on MB

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0333 $\pm$ 0.0577	0.0067 $\pm$ 0.0115	0.0111 $\pm$ 0.0192	0.0056 $\pm$ 0.0098	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0126 $\pm$ 0.0218	0.0067 $\pm$ 0.0115	0.0087 $\pm$ 0.0151	0.0044 $\pm$ 0.0076	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0177 $\pm$ 0.0215	0.0100 $\pm$ 0.0100	0.0102 $\pm$ 0.0089	0.0052 $\pm$ 0.0045	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0282 $\pm$ 0.0237	0.0167 $\pm$ 0.0058	0.0192 $\pm$ 0.0090	0.0097 $\pm$ 0.0046	0.0667 $\pm$ 0.1155	0.0333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics	0.5308 $\pm$ 0.0327	0.5033 $\pm$ 0.0551	0.5164 $\pm$ 0.0440	0.3488 $\pm$ 0.0394	0.8667 $\pm$ 0.1155	0.8000 $\pm$ 0.1000
L3-FT(CRISPR) + Omics + KG	0.5086 $\pm$ 0.0633	0.5000 $\pm$ 0.0608	0.5043 $\pm$ 0.0620	0.3387 $\pm$ 0.0569	0.7333 $\pm$ 0.2309	0.7000 $\pm$ 0.2000
G-Retriever + pre-GAT	<b>0.5546<math>\pm</math>0.0421</b>	0.3600 $\pm$ 0.0173	0.4365 $\pm$ 0.0257	0.2794 $\pm$ 0.0213	0.8000 $\pm$ 0.0000	0.7000 $\pm$ 0.0000
RoG	0.5530 $\pm$ 0.0211	<b>0.5400<math>\pm</math>0.0173</b>	<b>0.5464<math>\pm</math>0.0192</b>	<b>0.3761<math>\pm</math>0.0183</b>	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0000
SubgraphRAG	0.5545 $\pm$ 0.0119	0.3900 $\pm$ 0.1212	0.4502 $\pm$ 0.0737	0.2925 $\pm$ 0.0632	0.8667 $\pm$ 0.1155	0.9000 $\pm$ 0.0000
GNN-RAG	0.5306 $\pm$ 0.0267	0.4467 $\pm$ 0.1097	0.4776 $\pm$ 0.0606	0.3151 $\pm$ 0.0512	0.9333 $\pm$ 0.1155	<b>0.9667<math>\pm</math>0.0577</b>
GALAX	0.4897 $\pm$ 0.0304	0.4867 $\pm$ 0.0351	0.4882 $\pm$ 0.0328	0.3233 $\pm$ 0.0286	<b>1.0000<math>\pm</math>0.0000</b>	0.8333 $\pm$ 0.0577
GALAX (Qwen2.5-7B)	0.5300 $\pm$ 0.0346	0.5300 $\pm$ 0.0346	0.5300 $\pm$ 0.0346	0.3611 $\pm$ 0.0325	0.8667 $\pm$ 0.1155	0.9333 $\pm$ 0.0577

**Table 31:** Model performance on ALL

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0082	0.0050	0.0062	0.0031	0.0000	0.0500
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0333 $\pm$ 0.0577	0.0033 $\pm$ 0.0058	0.0061 $\pm$ 0.0105	0.0031 $\pm$ 0.0053	0.0667 $\pm$ 0.1155	0.0333 $\pm$ 0.0577
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0015 $\pm$ 0.0026	0.0033 $\pm$ 0.0058	0.0021 $\pm$ 0.0036	0.0010 $\pm$ 0.0018	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0607 $\pm$ 0.0528	0.0133 $\pm$ 0.0115	0.0219 $\pm$ 0.0189	0.0111 $\pm$ 0.0096	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5487 $\pm$ 0.0561	0.5400 $\pm$ 0.0624	0.5443 $\pm$ 0.0592	0.3754 $\pm$ 0.0569	<b>1.0000<math>\pm</math>0.0000</b>	0.9333 $\pm$ 0.1155
L3-FT(CRISPR) + Omics + KG	0.6010 $\pm$ 0.0608	0.5600 $\pm$ 0.1100	0.5786 $\pm$ 0.0868	0.4106 $\pm$ 0.0872	0.9333 $\pm$ 0.1155	0.9333 $\pm$ 0.1155
G-Retriever + pre-GAT	0.6285 $\pm$ 0.0043	0.5300 $\pm$ 0.0000	0.5750 $\pm$ 0.0018	0.4036 $\pm$ 0.0018	0.8000 $\pm$ 0.0000	0.8000 $\pm$ 0.0000
RoG	0.5596 $\pm$ 0.0699	0.5133 $\pm$ 0.1501	0.5326 $\pm$ 0.1168	0.3684 $\pm$ 0.1042	0.9333 $\pm$ 0.1155	0.9667 $\pm$ 0.0577
SubgraphRAG	0.6000 $\pm$ 0.0000	0.6000 $\pm$ 0.0000	0.6000 $\pm$ 0.0000	0.4286 $\pm$ 0.0000	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GNN-RAG	0.5682 $\pm$ 0.0858	0.5200 $\pm$ 0.1609	0.5402 $\pm$ 0.1293	0.3770 $\pm$ 0.1186	0.9333 $\pm$ 0.1155	0.9667 $\pm$ 0.0577
GALAX	<b>0.6860<math>\pm</math>0.0246</b>	<b>0.6700<math>\pm</math>0.0265</b>	<b>0.6779<math>\pm</math>0.0255</b>	<b>0.5131<math>\pm</math>0.0295</b>	<b>1.0000<math>\pm</math>0.0000</b>	0.9333 $\pm$ 0.0577
GALAX (Qwen2.5-7B)	0.6000 $\pm$ 0.0520	0.6000 $\pm$ 0.0520	0.6000 $\pm$ 0.0520	0.4299 $\pm$ 0.0542	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>

**Table 32:** Model performance on LGG

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0053 $\pm$ 0.0092	0.0033 $\pm$ 0.0058	0.0041 $\pm$ 0.0071	0.0021 $\pm$ 0.0036	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics	0.0070 $\pm$ 0.0121	0.0133 $\pm$ 0.0231	0.0092 $\pm$ 0.0159	0.0046 $\pm$ 0.0080	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.5000 $\pm$ 0.0500	0.4367 $\pm$ 0.0702	0.4649 $\pm$ 0.0545	0.3040 $\pm$ 0.0457	<b>1.0000<math>\pm</math>0.0000</b>	0.9667 $\pm$ 0.0577
L3-FT(CRISPR) + Omics + KG	0.5267 $\pm$ 0.0250	0.5067 $\pm$ 0.0115	0.5163 $\pm$ 0.0141	0.3480 $\pm$ 0.0128	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.1000
G-Retriever + pre-GAT	0.4846 $\pm$ 0.0537	0.4300 $\pm$ 0.0173	0.4540 $\pm$ 0.0126	0.2937 $\pm$ 0.0106	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0000
RoG	<b>0.5667<math>\pm</math>0.0462</b>	<b>0.5667<math>\pm</math>0.0462</b>	<b>0.5667<math>\pm</math>0.0462</b>	<b>0.3963<math>\pm</math>0.0458</b>	<b>1.0000<math>\pm</math>0.0000</b>	0.9667 $\pm$ 0.0577
SubgraphRAG	0.5400 $\pm$ 0.0000	0.5400 $\pm$ 0.0000	0.5400 $\pm$ 0.0000	0.3699 $\pm$ 0.0000	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GNN-RAG	0.5294 $\pm$ 0.0000	0.2700 $\pm$ 0.0000	0.3576 $\pm$ 0.0000	0.2177 $\pm$ 0.0000	0.8000 $\pm$ 0.0000	0.8000 $\pm$ 0.0000
GALAX	0.5433 $\pm$ 0.0252	0.5433 $\pm$ 0.0252	0.5433 $\pm$ 0.0252	0.3733 $\pm$ 0.0238	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0000
GALAX (Qwen2.5-7B)	0.5501 $\pm$ 0.0576	0.5200 $\pm$ 0.0693	0.5345 $\pm$ 0.0639	0.3664 $\pm$ 0.0581	0.8667 $\pm$ 0.1155	0.8667 $\pm$ 0.0577

**Table 33:** Model performance on NB

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0085 $\pm$ 0.0148	0.0167 $\pm$ 0.0289	0.0113 $\pm$ 0.0196	0.0057 $\pm$ 0.0100	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.5597 $\pm$ 0.0479	0.5233 $\pm$ 0.0651	0.5407 $\pm$ 0.0559	0.3718 $\pm$ 0.0523	0.8667 $\pm$ 0.2309	0.8667 $\pm$ 0.1528
L3-FT(CRISPR) + Omics + KG	0.5507 $\pm$ 0.0704	0.5067 $\pm$ 0.0723	0.5275 $\pm$ 0.0703	0.3603 $\pm$ 0.0667	0.9333 $\pm$ 0.1155	0.8667 $\pm$ 0.0577
G-Retriever + pre-GAT	0.4322 $\pm$ 0.0444	0.3933 $\pm$ 0.0404	0.4119 $\pm$ 0.0423	0.2599 $\pm$ 0.0341	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9333<math>\pm</math>0.0577</b>
RoG	<b>0.5997<math>\pm</math>0.0551</b>	0.4133 $\pm$ 0.0924	0.4817 $\pm$ 0.0400	0.3179 $\pm$ 0.0353	<b>1.0000<math>\pm</math>0.0000</b>	0.8667 $\pm$ 0.1155
SubgraphRAG	0.5679 $\pm$ 0.0551	0.4667 $\pm$ 0.0924	0.5048 $\pm$ 0.0400	0.3383 $\pm$ 0.0353	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9333<math>\pm</math>0.1155</b>
GNN-RAG	0.5810 $\pm$ 0.0165	<b>0.5733<math>\pm</math>0.0231</b>	<b>0.5771<math>\pm</math>0.0198</b>	<b>0.4058<math>\pm</math>0.0197</b>	0.7333 $\pm$ 0.2309	0.7333 $\pm$ 0.0577
GALAX	0.5533 $\pm$ 0.0115	0.5533 $\pm$ 0.0115	0.5533 $\pm$ 0.0115	0.3825 $\pm$ 0.0110	0.8667 $\pm$ 0.1155	0.8667 $\pm$ 0.0577
GALAX (Qwen2.5-7B)	0.5733 $\pm$ 0.0462	<b>0.5733<math>\pm</math>0.0462</b>	0.5733 $\pm$ 0.0462	0.4028 $\pm$ 0.0446	0.8000 $\pm$ 0.0000	0.8667 $\pm$ 0.0577

**Table 34:** Model performance on MESO

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0070	0.0050	0.0058	0.0029	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0190 $\pm$ 0.0201	0.0167 $\pm$ 0.0208	0.0133 $\pm$ 0.0121	0.0067 $\pm$ 0.0061	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0192 $\pm$ 0.0228	0.0100 $\pm$ 0.0000	0.0104 $\pm$ 0.0052	0.0053 $\pm$ 0.0026	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics	0.5182 $\pm$ 0.0806	0.4833 $\pm$ 0.0874	0.4995 $\pm$ 0.0808	0.3354 $\pm$ 0.0720	0.9333 $\pm$ 0.1155	0.8667 $\pm$ 0.1528
L3-FT(CRISPR) + Omics + KG	0.5417 $\pm$ 0.0419	0.5233 $\pm$ 0.0416	0.5322 $\pm$ 0.0402	0.3633 $\pm$ 0.0374	0.8667 $\pm$ 0.1155	0.9000 $\pm$ 0.0000
G-Retriever + pre-GAT	0.4721 $\pm$ 0.0072	0.3733 $\pm$ 0.0462	0.4161 $\pm$ 0.0324	0.2630 $\pm$ 0.0255	<b>1.0000<math>\pm</math>0.0000</b>	0.8000 $\pm$ 0.0000
RoG	0.4476 $\pm$ 0.0755	0.3900 $\pm$ 0.1375	0.4137 $\pm$ 0.1135	0.2651 $\pm$ 0.0893	0.8667 $\pm$ 0.2309	0.8000 $\pm$ 0.0000
SubgraphRAG	0.4421 $\pm$ 0.0000	0.4200 $\pm$ 0.0000	0.4308 $\pm$ 0.0000	0.2745 $\pm$ 0.0000	0.6000 $\pm$ 0.0000	0.8000 $\pm$ 0.0000
GNN-RAG	0.3750 $\pm$ 0.0000	0.2400 $\pm$ 0.0000	0.2927 $\pm$ 0.0000	0.1714 $\pm$ 0.0000	<b>1.0000<math>\pm</math>0.0000</b>	0.8000 $\pm$ 0.0000
GALAX	<b>0.5933<math>\pm</math>0.0231</b>	<b>0.5933<math>\pm</math>0.0231</b>	<b>0.5933<math>\pm</math>0.0231</b>	<b>0.4221<math>\pm</math>0.0236</b>	<b>1.0000<math>\pm</math>0.0000</b>	0.8000 $\pm$ 0.0000
GALAX (Qwen2.5-7B)	0.5621 $\pm$ 0.0310	0.5533 $\pm$ 0.0462	0.5576 $\pm$ 0.0388	0.3872 $\pm$ 0.0367	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9333<math>\pm</math>0.0577</b>

**Table 35:** Model performance on LIHC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0100 $\pm$ 0.0000	0.0100 $\pm$ 0.0000	0.0100 $\pm$ 0.0000	0.0050 $\pm$ 0.0000	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0031 $\pm$ 0.0027	0.0067 $\pm$ 0.0058	0.0043 $\pm$ 0.0037	0.0021 $\pm$ 0.0019	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.4239 $\pm$ 0.0343	0.3967 $\pm$ 0.0513	0.4096 $\pm$ 0.0435	0.2582 $\pm$ 0.0341	0.7333 $\pm$ 0.1155	0.7000 $\pm$ 0.1732
L3-FT(CRISPR) + Omics + KG	0.4552 $\pm$ 0.0568	0.4467 $\pm$ 0.0635	0.4509 $\pm$ 0.0600	0.2924 $\pm$ 0.0512	0.8000 $\pm$ 0.0000	0.7667 $\pm$ 0.0577
G-Retriever + pre-GAT	0.3598 $\pm$ 0.0387	0.3000 $\pm$ 0.0520	0.3270 $\pm$ 0.0467	0.1961 $\pm$ 0.0340	0.8000 $\pm$ 0.0000	0.7667 $\pm$ 0.1155
RoG	0.4205 $\pm$ 0.1173	0.3733 $\pm$ 0.1856	0.3903 $\pm$ 0.1596	0.2507 $\pm$ 0.1247	0.9333 $\pm$ 0.1155	0.8000 $\pm$ 0.1732
SubgraphRAG	0.4738 $\pm$ 0.1320	0.4267 $\pm$ 0.2136	0.4436 $\pm$ 0.1843	0.2964 $\pm$ 0.1437	0.9333 $\pm$ 0.1155	0.8000 $\pm$ 0.1732
GNN-RAG	0.4233 $\pm$ 0.0289	0.4233 $\pm$ 0.0289	0.4233 $\pm$ 0.0289	0.2688 $\pm$ 0.0230	<b>1.0000<math>\pm</math>0.0000</b>	0.8333 $\pm$ 0.0577
GALAX	<b>0.4900<math>\pm</math>0.0100</b>	<b>0.4900<math>\pm</math>0.0100</b>	<b>0.4900<math>\pm</math>0.0100</b>	<b>0.3245<math>\pm</math>0.0088</b>	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9000<math>\pm</math>0.0000</b>
GALAX (Qwen2.5-7B)	0.4362 $\pm$ 0.0033	0.4333 $\pm$ 0.0058	0.4348 $\pm$ 0.0045	0.2778 $\pm$ 0.0037	0.8667 $\pm$ 0.1155	0.8667 $\pm$ 0.0577

**Table 36:** Model performance on LAML

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0164	0.0100	0.0124	0.0062	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.1576 $\pm$ 0.1412	0.0200 $\pm$ 0.0173	0.0354 $\pm$ 0.0307	0.0182 $\pm$ 0.0158	0.1333 $\pm$ 0.2309	0.2000 $\pm$ 0.1732
L3-FT(Med) + Omics	0.1136 $\pm$ 0.0630	0.0400 $\pm$ 0.0173	0.0495 $\pm$ 0.0008	0.0254 $\pm$ 0.0004	0.2000 $\pm$ 0.3464	0.1333 $\pm$ 0.1528
L3-FT(Med) + Omics + KG	0.1085 $\pm$ 0.0867	0.0267 $\pm$ 0.0058	0.0377 $\pm$ 0.0221	0.0193 $\pm$ 0.0114	0.2000 $\pm$ 0.3464	0.2000 $\pm$ 0.1732
L3-FT(CRISPR) + Omics	0.4758 $\pm$ 0.0251	0.3733 $\pm$ 0.0115	0.4180 $\pm$ 0.0082	0.2642 $\pm$ 0.0066	0.8667 $\pm$ 0.2309	0.9000 $\pm$ 0.1000
L3-FT(CRISPR) + Omics + KG	<b>0.4988<math>\pm</math>0.0700</b>	0.4500 $\pm$ 0.0600	<b>0.4721<math>\pm</math>0.0572</b>	<b>0.3102<math>\pm</math>0.0481</b>	0.8667 $\pm$ 0.2309	0.8333 $\pm$ 0.1155
G-Retriever + pre-GAT	0.3719 $\pm$ 0.0373	0.2400 $\pm$ 0.0000	0.2912 $\pm$ 0.0119	0.1705 $\pm$ 0.0081	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0000
RoG	0.4228 $\pm$ 0.0149	0.3467 $\pm$ 0.1097	0.3725 $\pm$ 0.0685	0.2303 $\pm$ 0.0506	<b>1.0000<math>\pm</math>0.0000</b>	0.8667 $\pm$ 0.1155
SubgraphRAG	0.4141 $\pm$ 0.0000	0.4100 $\pm$ 0.0000	0.4121 $\pm$ 0.0000	0.2595 $\pm$ 0.0000	<b>1.0000<math>\pm</math>0.0000</b>	0.8000 $\pm$ 0.0000
GNN-RAG	0.3972 $\pm$ 0.0934	0.3200 $\pm$ 0.1559	0.3511 $\pm$ 0.1311	0.2183 $\pm$ 0.1015	<b>1.0000<math>\pm</math>0.0000</b>	0.8333 $\pm$ 0.0577
GALAX	0.4730 $\pm$ 0.0901	0.4333 $\pm$ 0.0306	0.4510 $\pm$ 0.0554	0.2923 $\pm$ 0.0467	0.9333 $\pm$ 0.1155	0.8667 $\pm$ 0.2309
GALAX (Qwen2.5-7B)	0.4596 $\pm$ 0.0437	<b>0.4567<math>\pm</math>0.0462</b>	0.4581 $\pm$ 0.0449	0.2979 $\pm$ 0.0385	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>

**Table 37:** Model performance on DLBC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0220 $\pm$ 0.0249	0.0367 $\pm$ 0.0404	0.0275 $\pm$ 0.0308	0.0141 $\pm$ 0.0159	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0375 $\pm$ 0.0385	0.0067 $\pm$ 0.0058	0.0111 $\pm$ 0.0097	0.0056 $\pm$ 0.0049	0.0667 $\pm$ 0.1155	0.0333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics	0.3985 $\pm$ 0.0669	0.3900 $\pm$ 0.0721	0.3942 $\pm$ 0.0695	0.2470 $\pm$ 0.0530	0.9333 $\pm$ 0.1155	0.8333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics + KG	0.4200 $\pm$ 0.0230	0.4100 $\pm$ 0.0173	0.4149 $\pm$ 0.0200	0.2619 $\pm$ 0.0158	0.8667 $\pm$ 0.1155	0.7000 $\pm$ 0.1000
G-Retriever + pre-GAT	0.3196 $\pm$ 0.0307	0.2767 $\pm$ 0.0058	0.2963 $\pm$ 0.0168	0.1740 $\pm$ 0.0115	<b>1.0000<math>\pm</math>0.0000</b>	0.8000 $\pm$ 0.0000
RoG	0.4424 $\pm$ 0.0239	0.4367 $\pm$ 0.0289	0.4395 $\pm$ 0.0264	0.2819 $\pm$ 0.0219	0.8667 $\pm$ 0.1155	0.7667 $\pm$ 0.1155
SubgraphRAG	0.4286 $\pm$ 0.0000	0.4200 $\pm$ 0.0000	0.4242 $\pm$ 0.0000	0.2692 $\pm$ 0.0000	0.8000 $\pm$ 0.0000	0.7000 $\pm$ 0.0000
GNN-RAG	0.4433 $\pm$ 0.0231	0.4433 $\pm$ 0.0231	0.4433 $\pm$ 0.0231	0.2850 $\pm$ 0.0192	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9000<math>\pm</math>0.0000</b>
GALAX	0.4400 $\pm$ 0.0000	0.4400 $\pm$ 0.0000	0.4400 $\pm$ 0.0000	0.2821 $\pm$ 0.0000	0.8667 $\pm$ 0.1155	0.6333 $\pm$ 0.2309
GALAX (Qwen2.5-7B)	<b>0.4764<math>\pm</math>0.0055</b>	<b>0.4700<math>\pm</math>0.0000</b>	<b>0.4732<math>\pm</math>0.0027</b>	<b>0.3099<math>\pm</math>0.0023</b>	0.8000 $\pm$ 0.0000	<b>0.9000<math>\pm</math>0.0000</b>

**Table 38:** Model performance on MM

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0275 $\pm$ 0.0296	0.0167 $\pm$ 0.0208	0.0156 $\pm$ 0.0149	0.0079 $\pm$ 0.0075	0.0667 $\pm$ 0.1155	0.0333 $\pm$ 0.0577
L3-FT(Med) + Omics + KG	0.1091 $\pm$ 0.0630	0.0200 $\pm$ 0.0000	0.0329 $\pm$ 0.0027	0.0167 $\pm$ 0.0014	0.1333 $\pm$ 0.1155	0.1333 $\pm$ 0.1155
L3-FT(CRISPR) + Omics	0.4001 $\pm$ 0.0022	0.3800 $\pm$ 0.0265	0.3894 $\pm$ 0.0134	0.2418 $\pm$ 0.0103	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics + KG	0.4753 $\pm$ 0.0613	0.4367 $\pm$ 0.0306	0.4549 $\pm$ 0.0447	0.2952 $\pm$ 0.0377	0.8000 $\pm$ 0.0000	0.7333 $\pm$ 0.0577
G-Retriever + pre-GAT	0.4514 $\pm$ 0.0890	0.3667 $\pm$ 0.0808	0.4046 $\pm$ 0.0850	0.2561 $\pm$ 0.0690	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.0000
RoG	0.4900 $\pm$ 0.0615	0.3833 $\pm$ 0.1550	0.4246 $\pm$ 0.1182	0.2745 $\pm$ 0.0997	0.9333 $\pm$ 0.1155	0.8667 $\pm$ 0.0577
SubgraphRAG	0.4515 $\pm$ 0.0122	0.3033 $\pm$ 0.0289	0.3620 $\pm$ 0.0175	0.2211 $\pm$ 0.0130	0.9333 $\pm$ 0.1155	0.8667 $\pm$ 0.0577
GNN-RAG	0.5166 $\pm$ 0.0376	0.5133 $\pm$ 0.0404	0.5150 $\pm$ 0.0390	0.3474 $\pm$ 0.0359	<b>1.0000<math>\pm</math>0.0000</b>	0.8333 $\pm$ 0.0577
GALAX	<b>0.5858<math>\pm</math>0.0564</b>	<b>0.5700<math>\pm</math>0.0819</b>	<b>0.5775<math>\pm</math>0.0695</b>	<b>0.4082<math>\pm</math>0.0680</b>	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9333<math>\pm</math>0.0577</b>
GALAX (Qwen2.5-7B)	0.5000 $\pm$ 0.0917	0.5000 $\pm$ 0.0917	0.5000 $\pm$ 0.0917	0.3367 $\pm$ 0.0834	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.1000

**Table 39:** Model performance on KIRC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	0.5458 $\pm$ 0.0450	0.5333 $\pm$ 0.0513	0.5394 $\pm$ 0.0478	0.3703 $\pm$ 0.0452	0.8667 $\pm$ 0.2309	0.8667 $\pm$ 0.1155
L3-FT(CRISPR) + Omics + KG	0.4897 $\pm$ 0.0182	0.4800 $\pm$ 0.0265	0.4847 $\pm$ 0.0219	0.3201 $\pm$ 0.0192	0.8667 $\pm$ 0.1155	0.9000 $\pm$ 0.0000
G-Retriever + pre-GAT	0.4160 $\pm$ 0.0315	0.3733 $\pm$ 0.0058	0.3933 $\pm$ 0.0171	0.2449 $\pm$ 0.0133	0.6000 $\pm$ 0.0000	0.8000 $\pm$ 0.0000
RoG	0.6207 $\pm$ 0.0617	<b>0.5867<math>\pm</math>0.0751</b>	<b>0.6031<math>\pm</math>0.0689</b>	<b>0.4340<math>\pm</math>0.0687</b>	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
SubgraphRAG	<b>0.6412<math>\pm</math>0.0261</b>	0.5300 $\pm$ 0.1732	0.5714 $\pm$ 0.1237	0.4067 $\pm$ 0.1160	<b>1.0000<math>\pm</math>0.0000</b>	0.9333 $\pm$ 0.1155
GNN-RAG	0.5898 $\pm$ 0.0350	0.5733 $\pm$ 0.0635	0.5812 $\pm$ 0.0499	0.4108 $\pm$ 0.0480	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GALAX	0.5797 $\pm$ 0.0251	0.5700 $\pm$ 0.0265	0.5748 $\pm$ 0.0254	0.4036 $\pm$ 0.0247	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.1732
GALAX (Qwen2.5-7B)	0.5786 $\pm$ 0.1114	0.5700 $\pm$ 0.1044	0.5742 $\pm$ 0.1076	0.4079 $\pm$ 0.1019	0.8667 $\pm$ 0.1155	0.9333 $\pm$ 0.0577

**Table 40:** Model performance on THCA

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.1500 $\pm$ 0.1500	0.0200 $\pm$ 0.0173	0.0348 $\pm$ 0.0303	0.0179 $\pm$ 0.0155	0.1333 $\pm$ 0.2309	0.1000 $\pm$ 0.1732
L3-FT(Med) + Omics	0.0439 $\pm$ 0.0532	0.0167 $\pm$ 0.0058	0.0200 $\pm$ 0.0119	0.0101 $\pm$ 0.0061	0.0667 $\pm$ 0.1155	0.0333 $\pm$ 0.0577
L3-FT(Med) + Omics + KG	0.0156 $\pm$ 0.0145	0.0300 $\pm$ 0.0265	0.0205 $\pm$ 0.0187	0.0104 $\pm$ 0.0095	0.0000 $\pm$ 0.0000	0.0333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics	0.5495 $\pm$ 0.0165	0.5367 $\pm$ 0.0153	0.5430 $\pm$ 0.0158	0.3728 $\pm$ 0.0148	0.8667 $\pm$ 0.1155	0.8333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics + KG	0.5818 $\pm$ 0.0490	0.5333 $\pm$ 0.0416	0.5556 $\pm$ 0.0365	0.3853 $\pm$ 0.0353	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.1000
G-Retriever + pre-GAT	0.6237 $\pm$ 0.0849	0.3833 $\pm$ 0.0231	0.4718 $\pm$ 0.0097	0.3088 $\pm$ 0.0083	0.5333 $\pm$ 0.1155	0.7333 $\pm$ 0.1155
RoG	0.5696 $\pm$ 0.0350	0.5367 $\pm$ 0.0635	0.5524 $\pm$ 0.0499	0.3827 $\pm$ 0.0486	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
SubgraphRAG	0.5495 $\pm$ 0.0000	0.5000 $\pm$ 0.0000	0.5236 $\pm$ 0.0000	0.3546 $\pm$ 0.0000	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GNN-RAG	<b>0.6465<math>\pm</math>0.0350</b>	0.4400 $\pm$ 0.1386	0.5127 $\pm$ 0.0782	0.3473 $\pm$ 0.0731	0.9333 $\pm$ 0.1155	0.8333 $\pm$ 0.0577
GALAX	0.6076 $\pm$ 0.0331	<b>0.6033<math>\pm</math>0.0289</b>	<b>0.6054<math>\pm</math>0.0308</b>	<b>0.4346<math>\pm</math>0.0313</b>	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GALAX (Qwen2.5-7B)	0.5790 $\pm$ 0.0350	0.4967 $\pm$ 0.0751	0.5312 $\pm$ 0.0356	0.3622 $\pm$ 0.0333	0.8000 $\pm$ 0.2000	0.8333 $\pm$ 0.0577

**Table 41:** Model performance on BLCA

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0127 $\pm$ 0.0111	0.0267 $\pm$ 0.0231	0.0172 $\pm$ 0.0149	0.0087 $\pm$ 0.0076	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0222 $\pm$ 0.0385	0.0033 $\pm$ 0.0058	0.0058 $\pm$ 0.0100	0.0029 $\pm$ 0.0051	0.0000 $\pm$ 0.0000	0.0333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics	0.5538 $\pm$ 0.0498	0.5000 $\pm$ 0.0200	0.5250 $\pm$ 0.0301	0.3563 $\pm$ 0.0279	0.7333 $\pm$ 0.1155	0.8000 $\pm$ 0.0000
L3-FT(CRISPR) + Omics + KG	0.5617 $\pm$ 0.0584	0.5367 $\pm$ 0.0231	0.5484 $\pm$ 0.0376	0.3784 $\pm$ 0.0356	0.8667 $\pm$ 0.2309	0.8667 $\pm$ 0.0577
G-Retriever + pre-GAT	0.5025 $\pm$ 0.0481	0.4733 $\pm$ 0.0058	0.4870 $\pm$ 0.0253	0.3221 $\pm$ 0.0223	0.6000 $\pm$ 0.0000	0.6333 $\pm$ 0.0577
RoG	0.5633 $\pm$ 0.0231	0.5633 $\pm$ 0.0231	0.5633 $\pm$ 0.0231	0.3924 $\pm$ 0.0226	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9667<math>\pm</math>0.0577</b>
SubgraphRAG	0.5457 $\pm$ 0.0075	0.4633 $\pm$ 0.1501	0.4922 $\pm$ 0.1001	0.3302 $\pm$ 0.0851	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.1732
GNN-RAG	0.5270 $\pm$ 0.1006	0.4900 $\pm$ 0.1646	0.5056 $\pm$ 0.1376	0.3456 $\pm$ 0.1177	0.6667 $\pm$ 0.5774	0.6000 $\pm$ 0.5196
GALAX	0.5247 $\pm$ 0.0081	0.5033 $\pm$ 0.0289	0.5133 $\pm$ 0.0115	0.3453 $\pm$ 0.0104	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.1000
GALAX (Qwen2.5-7B)	<b>0.5852<math>\pm</math>0.0091</b>	<b>0.5733<math>\pm</math>0.0115</b>	<b>0.5791<math>\pm</math>0.0015</b>	<b>0.4076<math>\pm</math>0.0015</b>	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.0000

**Table 42:** Model performance on UCEC

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0000 $\pm$ 0.0000					
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(CRISPR) + Omics	<b>0.5293<math>\pm</math>0.0394</b>	0.5133 $\pm$ 0.0379	0.5211 $\pm$ 0.0379	0.3530 $\pm$ 0.0344	0.7333 $\pm$ 0.1155	0.7000 $\pm$ 0.1000
L3-FT(CRISPR) + Omics + KG	0.4672 $\pm$ 0.0472	0.4533 $\pm$ 0.0231	0.4600 $\pm$ 0.0346	0.2991 $\pm$ 0.0296	0.8000 $\pm$ 0.2000	0.8000 $\pm$ 0.1000
G-Retriever + pre-GAT	0.4517 $\pm$ 0.0318	0.4000 $\pm$ 0.0173	0.4242 $\pm$ 0.0237	0.2694 $\pm$ 0.0193	<b>1.0000<math>\pm</math>0.0000</b>	<b>0.9000<math>\pm</math>0.0000</b>
RoG	0.4508 $\pm$ 0.0506	0.4233 $\pm$ 0.0982	0.4355 $\pm$ 0.0772	0.2803 $\pm$ 0.0614	0.6667 $\pm$ 0.5774	0.4667 $\pm$ 0.4041
SubgraphRAG	0.4800 $\pm$ 0.0000	0.4800 $\pm$ 0.0000	0.4800 $\pm$ 0.0000	0.3158 $\pm$ 0.0000	<b>1.0000<math>\pm</math>0.0000</b>	0.7000 $\pm$ 0.0000
GNN-RAG	0.4777 $\pm$ 0.0739	0.4433 $\pm$ 0.1159	0.4589 $\pm$ 0.0975	0.3011 $\pm$ 0.0795	0.6000 $\pm$ 0.5292	0.4667 $\pm$ 0.4163
GALAX	0.5280 $\pm$ 0.0555	<b>0.5233<math>\pm</math>0.0635</b>	<b>0.5256<math>\pm</math>0.0596</b>	<b>0.3579<math>\pm</math>0.0536</b>	0.8000 $\pm$ 0.0000	<b>0.9000<math>\pm</math>0.0000</b>
GALAX (Qwen2.5-7B)	0.4848 $\pm$ 0.0391	0.5033 $\pm$ 0.0231	0.4938 $\pm$ 0.0313	0.3282 $\pm$ 0.0280	0.9333 $\pm$ 0.1155	0.8333 $\pm$ 0.0577

**Table 43:** Model performance on PRAD

Model	Precision $\uparrow$	Recall $\uparrow$	F1 $\uparrow$	Jaccard $\uparrow$	Hit@5 $\uparrow$	Hit@10 $\uparrow$
M2T	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
GAT	0.0000 $\pm$ 0.0000					
L3 + Omics	0.0111 $\pm$ 0.0192	0.0033 $\pm$ 0.0058	0.0051 $\pm$ 0.0089	0.0026 $\pm$ 0.0045	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3 + Omics + KG	0.0000 $\pm$ 0.0000					
L3-FT(Med) + Omics	0.0145 $\pm$ 0.0251	0.0033 $\pm$ 0.0058	0.0054 $\pm$ 0.0094	0.0027 $\pm$ 0.0047	0.0000 $\pm$ 0.0000	0.0000 $\pm$ 0.0000
L3-FT(Med) + Omics + KG	0.0176 $\pm$ 0.0261	0.0067 $\pm$ 0.0058	0.0078 $\pm$ 0.0083	0.0039 $\pm$ 0.0042	0.0000 $\pm$ 0.0000	0.0333 $\pm$ 0.0577
L3-FT(CRISPR) + Omics	0.6021 $\pm$ 0.0558	0.5600 $\pm$ 0.0400	0.5796 $\pm$ 0.0401	0.4088 $\pm$ 0.0392	<b>1.0000<math>\pm</math>0.0000</b>	0.9333 $\pm$ 0.1155
L3-FT(CRISPR) + Omics + KG	0.6032 $\pm$ 0.0191	0.5867 $\pm$ 0.0115	0.5947 $\pm$ 0.0102	0.4232 $\pm$ 0.0103	0.9333 $\pm$ 0.1155	0.9000 $\pm$ 0.1000
G-Retriever + pre-GAT	0.4749 $\pm$ 0.0041	0.4433 $\pm$ 0.0231	0.4584 $\pm$ 0.0141	0.2974 $\pm$ 0.0120	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
RoG	0.5833 $\pm$ 0.0404	0.5833 $\pm$ 0.0404	0.5833 $\pm$ 0.0404	0.4125 $\pm$ 0.0410	<b>1.0000<math>\pm</math>0.0000</b>	0.9333 $\pm$ 0.0577
SubgraphRAG	0.5931 $\pm$ 0.1258	0.4500 $\pm$ 0.0964	0.5042 $\pm$ 0.0691	0.3389 $\pm$ 0.0607	0.6667 $\pm$ 0.5774	0.6333 $\pm$ 0.5508
GNN-RAG	0.6140 $\pm$ 0.0138	0.6100 $\pm$ 0.0173	0.6120 $\pm$ 0.0156	0.4411 $\pm$ 0.0163	<b>1.0000<math>\pm</math>0.0000</b>	<b>1.0000<math>\pm</math>0.0000</b>
GALAX	0.5152 $\pm$ 0.1398	0.4733 $\pm$ 0.0751	0.4923 $\pm$ 0.1036	0.3308 $\pm$ 0.0934	0.9333 $\pm$ 0.1155	0.9667 $\pm$ 0.0577
GALAX (Qwen2.5-7B)	<b>0.6195<math>\pm</math>0.0292</b>	<b>0.6133<math>\pm</math>0.0289</b>	<b>0.6164<math>\pm</math>0.0290</b>	<b>0.4459<math>\pm</math>0.0300</b>	<b>1.0000<math>\pm</math>0.0000</b>	0.9000 $\pm$ 0.0000