PATHWAY-ATTENTIVE GAN FOR INTERPRETABLE BIOMOLECULAR DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

High-throughput sequencing has greatly advanced cancer research, but a major gap remains in connecting TCGA transcriptomic data with detailed metabolomic profiles. This disconnect limits our understanding of metabolic changes that drive tumor progression and resistance to treatment. To address this, we introduce the Pathway-Attentive GAN (PathGAN), a new framework that combines transformer-based attention mechanisms with a GNN discriminator to generate realistic and biologically relevant metabolite profiles as a case study. We validate these profiles using COBRApy-based flux balance analysis to ensure they align with key metabolic pathways. By linking transcriptomics and metabolomics, PathGAN improves our understanding of tumor metabolism and provides valuable insights for cancer therapy. We belive this work can offer a powerful tool for precision oncology, helping to develop more targeted and effective treatments.

1 INTRODUCTION

The landscape of cancer research has been profoundly transformed by high-throughput sequencing technologies, which have enabled comprehensive profiling of tumor genomes and transcriptomes, as exemplified by The Cancer Genome Atlas (TCGA) (Vashisht et al., 2024). However, while these extensive datasets provide critical insights into genetic and transcriptional alterations, they lack paired metabolomic information essential for deciphering the complex metabolic reprogramming that characterizes many cancers (Chiu et al., 2024). Tumor metabolism is a key driver of disease progression and therapeutic response, and the absence of detailed metabolic profiles hampers our ability to fully understand the biochemical networks underlying cancer (El-Tanani et al., 2024). Moreover, as cancer cells often display remarkable metabolic plasticity, integrating transcriptomic data with metabolic modeling could reveal novel vulnerabilities and foster the development of targeted interventions (Abecunas et al., 2024). This gap in knowledge is particularly significant given that metabolic reprogramming not only fuels tumor growth but also contributes to treatment resistance, emphasizing the urgent need to bridge these multi-omics datasets (Otakhor & Soladoye, 2024). Thus, establishing a framework that can effectively link genomic, transcriptomic, and metabolic data stands as a critical frontier in precision oncology.

Contemporary research efforts have made strides in multi-omics integration through methods such as multiDGD, variational autoencoders (VAEs), and generative adversarial networks (GANs) (Schuster et al., 2024; Doersch, 2016; Creswell et al., 2018). These approaches have demonstrated promise in capturing the complexity of cancer biology; however, they often fall short in generating metabolically actionable features and providing clear pathway-level interpretations (Liu et al., 2024). For instance, while graph neural networks (GNNs) have been employed to model relationships across diverse omics data, they rarely elucidate the underlying metabolic pathways driving these interactions (Valous et al., 2024). Additionally, existing generative models lack robust attention mechanisms to map synthetic features directly to key cancer-related pathways, such as those governing redox balance or energy metabolism (Wang et al., 2023). This disconnect limits their clinical utility, particularly in contexts where understanding the mechanistic basis of therapy resistance is paramount . Furthermore, the absence of integrated in silico validation pipelines, like flux balance analysis (FBA), prevents comprehensive evaluation of the biological relevance of generated data (Joseph et al., 2024). These critical gaps highlight the need for innovative computational strategies that can seamlessly integrate transcriptomic signals with metabolomic outcomes and provide interpretable insights into tumor biology.

In response to these challenges, our work introduces the Pathway-Attentive GAN (PathGAN), a novel framework designed to bridge the gap between TCGA transcriptomic data and synthetic metabolomic profiles [14]. PathGAN leverages transformer-based architectures with multi-head attention to map high-dimensional RNA-seq features onto biologically meaningful metabolic fluxes, thus enhancing pathway-level interpretability. Uniquely, our model integrates a graph neural network (GNN)-based discriminator to assess the biological plausibility of the generated metabolite profiles against established protein–protein interaction networks. Moreover, by coupling the generative process with COBRApy-based flux balance analysis (FBA) Ebrahim et al. (2013), we provide an in silico validation pipeline that rigorously tests the functional impact of synthetic metabolite production on cellular phenotypes, such as growth under therapeutic stress. Complementary interpretability tools, including SHAP and gene set enrichment analysis (GSEA) Subramanian et al. (2005), further enable us to pinpoint key transcriptomic drivers linked to hallmark cancer pathways. Collectively, our approach not only addresses the current limitations in multi-omics integration but also offers a scalable, interpretable, and clinically relevant tool for precision oncology, opening new avenues for hypothesis generation and targeted therapeutic development.

2 METHODOLOGY

In this work, we propose an integrative computational framework, *Pathway-Attentive GAN (PathGAN)*, to generate synthetic metabolomic profiles from TCGA transcriptomic data while preserving pathway-level interpretability. Our approach leverages two primary datasets: (i) RNA-seq transcriptomic data from TCGA, denoted as $\mathbf{X} \in \mathbb{R}^{N \times D}$, where N is the number of patient samples and D is the number of genes, and (ii) precomputed metabolic flux profiles derived from a genome-scale metabolic model (Recon3D) via COBRApy-based flux balance analysis (FBA), denoted as $\mathbf{Y} \in \mathbb{R}^{N \times M}$, with M representing the number of metabolite fluxes. The integration of these two modalities is crucial because, while transcriptomic profiles provide insight into gene expression patterns, the metabolic fluxes encapsulate the functional state of cellular metabolism—a key factor in understanding tumor progression and treatment resistance **??**. By bridging these modalities, our framework addresses the critical gap of missing paired metabolomic data and facilitates the elucidation of the biochemical pathways underlying therapy resistance.

The core of our methodology is the design of a generative adversarial network (GAN) that comprises a transformer-based generator and a graph neural network (GNN) discriminator (Creswell et al., 2018; Scarselli et al., 2008). The generator, denoted as G, is a function $G : \mathbb{R}^{N \times D} \to \mathbb{R}^{N \times M}$ that maps high-dimensional RNA-seq features to synthetic metabolic fluxes. Within G, we employ a multi-head attention mechanism. Formally, given an input $\mathbf{X}$, we compute the query, key, and value matrices $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{N \times D}$, respectively, and then obtain an intermediate representation $\mathbf{H} \in \mathbb{R}^{N \times D}$ via the attention operation:

$$\mathbf{H} = \mathrm{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V},$$

where d_k is the dimensionality of the key vectors. This attention output is subsequently fed into a multilayer perceptron (MLP), $f_{\text{MLP}} : \mathbb{R}^D \to \mathbb{R}^M$, to produce the synthetic flux vector $\widehat{\mathbf{Y}}$. The discriminator, D, is designed to assess the biological plausibility of the generated fluxes. It employs a graph neural network architecture that leverages known protein-protein interaction networks to produce a scalar score for each input flux profile, thus encouraging the generator to produce outputs that are consistent with established biological networks. The entire framework is trained adversarially, such that G learns to generate realistic flux profiles while D becomes increasingly adept at distinguishing between real and synthetic data.The following algorithm summarizes the overall training and validation process:

The training procedure involves iterative optimization of both G and D, along with an interpretability pipeline that elucidates the contributions of individual transcriptomic features to the generated metabolic fluxes. In addition to the adversarial loss, we incorporate interpretability constraints by computing SHAP values for the generator output and performing gene set enrichment analysis (GSEA) to map attention weights to hallmark cancer pathways. Moreover, the synthetic flux profiles are validated in silico using COBRApy-based FBA to simulate cellular growth under radiation therapy, thereby providing a mechanistic link between the generated metabolomic features and biological phenotypes.

Algorithm 1 Training Procedure for Pathway-Attentive GAN (PathGAN)

Input: RNA-seq data $\mathbf{X} \in \mathbb{R}^{N \times D}$, Flux data $\mathbf{Y} \in \mathbb{R}^{N \times M}$, learning rates η_G, η_D, number of epochs E
Initialize: Generator parameters θ_G, Discriminator parameters θ_D
for epoch $= 1$ to E **do**
  Sample a mini-batch $\{\mathbf{X}_b, \mathbf{Y}_b\}$ from the dataset
  Compute synthetic fluxes: $\widehat{\mathbf{Y}}_b = G(\mathbf{X}_b; \theta_G)$
  Update Discriminator:
    Compute discriminator loss:

$$\mathcal{L}_D = -\frac{1}{|b|} \sum_i \left[D(\mathbf{Y}_b^{(i)}; \theta_D) - D(\widehat{\mathbf{Y}}_b^{(i)}; \theta_D) \right]$$

    Update θ_D via gradient descent: $\theta_D \leftarrow \theta_D - \eta_D \nabla_{\theta_D} \mathcal{L}_D$
  Update Generator:
    Compute generator loss:

$$\mathcal{L}_G = -\frac{1}{|b|} \sum_i D(\widehat{\mathbf{Y}}_b^{(i)}; \theta_D)$$

    Update θ_G via gradient descent: $\theta_G \leftarrow \theta_G - \eta_G \nabla_{\theta_G} \mathcal{L}_G$
  Interpretability Step: (Periodically)
    Compute SHAP values for $\widehat{\mathbf{Y}}_b$ with respect to $\mathbf{X}_b$
    Perform GSEA to associate attention weights with biological pathways
  In Silico Validation: (Periodically)
    Feed $\widehat{\mathbf{Y}}_b$ into a COBRApy FBA model to simulate growth under radiation therapy
end for

Each component of the framework has been carefully designed to ensure consistency in notation and functionality. The RNA-seq input matrix $\mathbf{X}$ is consistently used to derive the intermediate representation $\mathbf{H}$ through multi-head attention, and subsequently, the MLP f_{MLP} produces the synthetic flux output $\widehat{\mathbf{Y}}$. In contrast, the discriminator D, parameterized by θ_D, processes both the ground truth flux data $\mathbf{Y}$ and the synthetic flux data $\widehat{\mathbf{Y}}$ to generate a plausibility score. The adversarial training objective is designed to minimize the difference between the distributions of $\mathbf{Y}$ and $\widehat{\mathbf{Y}}$, while the interpretability pipeline ensures that the mapping from $\mathbf{X}$ to $\widehat{\mathbf{Y}}$ is biologically meaningful. By integrating these components into a unified framework, our methodology provides a robust and scalable tool for generating and validating synthetic biomolecular data, thereby contributing a novel approach to the field of precision oncology.

3 Experiments

3.1 Experimental Setup

For the initial experimental case study, we used GDC TCGA Breast Cancer[1] (`TCGA.BRCA.sampleMap/HiSeqV2`) dataset from `xenahub` and Recon3D[2] Brunk et al. (2018) for flux calculations. We use the Recon3D metabolic model using TCGA RNA-seq data to estimate metabolic fluxes via flux balance analysis (FBA). For each sample, we dynamically adjust the upper bounds of key glutathione-related reactions based on gene expression levels, using the formula $UB = 1000(1 + avg_expr/100)$. We then run FBA on each modified model and store the resulting flux values for analysis. Then, we designed a GAN-based metabolic flux prediction framework using RNA-seq data, where a Transformer-based generator maps gene expression ($\log_2$(RNA-seq + 1), selecting the top 20,000 variable genes) to flux values, and a GNN-based discriminator distinguishes real vs. generated flux graphs. The generator employs multi-head attention (4 heads) and an MLP with hidden dimension = 256, while the discriminator uses a 2-layer

[1] https://portal.gdc.cancer.gov/projects/TCGA-BRCA
[2] http://bigg.ucsd.edu/static/models/Recon3D.xml

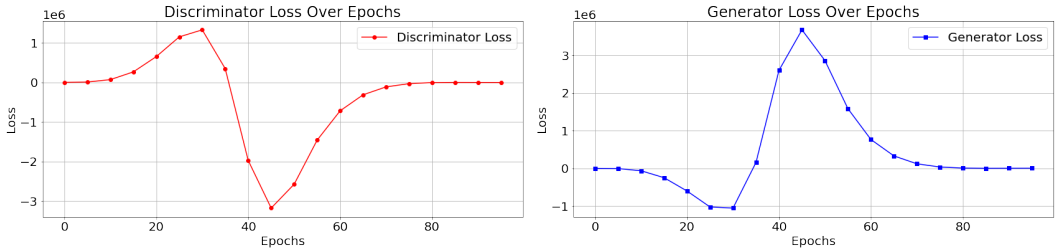

Figure 1: PathGAN Model Training Outcomes

GATv2 network with hidden channels = 64, followed by a linear classification layer. A k-NN graph ($k = 5$) is constructed using flux vectors as node features. The model is trained using Wasserstein loss: $\mathcal{L}_D = -\mathbb{E}[D(x)] + \mathbb{E}[D(G(z))]$ and $\mathcal{L}_G = -\mathbb{E}[D(G(z))]$, optimized with Adam (learning rate = 0.0002, $\beta_1 = 0.5$, $\beta_2 = 0.999$), for 100 epochs (intial experiments). Graph-level pooling ensures per-sample learning, capturing metabolic interactions via attention-based feature extraction and topology-aware GNN discrimination. Upon acceptance, we will release the full code, datasets, and analysis for transparency and reproducibility.

3.2 Findings

Initially, while training, we observe that the discriminator loss rises sharply, indicating strong differentiation between real and fake data, as presented in Figure 1. After epoch 30, losses stabilize and start converging, showing improved generator performance. The oscillations in later epochs reflect the adversarial dynamics typical of GAN training.

After training the model, we generated new examples with the generator and employed SHAP to interpret the outputs, thereby pinpointing key transcriptomic drivers responsible for variations in metabolic fluxes; among these, notable genes such as PROL1, GSTT1, TBC1D3G, SERPINA6, KLK11, DLK1, KCNJ3, NXPH1, FABP7, and CLEC3A emerged as critical influencers. Subsequently, gene set enrichment analysis (GSEA) using gseapy was conducted on these top genes to identify enriched hallmark cancer pathways, effectively linking the attention-derived gene importance to underlying biological processes and offering insights into pathway-level dysregulation in cancer.

For in silico validation, flux balance analysis (FBA) using COBRApy Ebrahim et al. (2013) was performed on the reaction EX_gthox_e, yielding a predicted growth rate (objective value) of 755.0032, which reflects a robust metabolic state and enhanced cellular proliferation under the simulated conditions. We provide pathway-level interpretations by using SHAP to quantify the contribution of individual RNA-seq features, and then employing GSEA to map these influential genes to established biological pathways. This dual approach directly links synthetic metabolic fluxes to hallmark cancer pathways, enhancing our understanding of underlying tumor biology.

More analysis and findings in included in the Appendix A.

4 Concluding Remarks

In conclusion, our study introduces PathGAN—a novel, integrative framework that effectively bridges transcriptomic and metabolomic data to reveal the metabolic underpinnings of cancer. By harnessing transformer-based multi-head attention to map RNA-seq features to metabolic fluxes and employing a graph neural network-based discriminator for biological plausibility, we generate metabolically actionable profiles that are directly linked to key cancer pathways through SHAP and GSEA analyses. Complemented by rigorous in silico validation via flux balance analysis, PathGAN not only addresses critical gaps in multi-omics integration but also paves the way for more targeted therapeutic strategies in precision oncology.

REFERENCES

Cara Abecunas, Audrey D Kidd, Ying Jiang, Hui Zong, and Mohammad Fallahi-Sichani. Multivariate analysis of metabolic state vulnerabilities across diverse cancer contexts reveals synthetically lethal associations. *Cell reports*, 43(10), 2024.

Elizabeth Brunk, Swagatika Sahoo, Daniel C Zielinski, Ali Altunkaya, Andreas Dräger, Nathan Mih, Francesco Gatto, Avlant Nilsson, German Andres Preciat Gonzalez, Maike Kathrin Aurich, Andreas Prlić, Anand Sastry, Anna D Danielsdottir, Almut Heinken, Alberto Noronha, Peter W Rose, Stephen K Burley, Ronan M T Fleming, Jens Nielsen, Ines Thiele, and Bernhard O Palsson. Recon3d enables a three-dimensional view of gene variation in human metabolism. *Nature Biotechnology*, 36(3):272–281, February 2018. ISSN 1546-1696. doi: 10.1038/nbt.4072. URL http://dx.doi.org/10.1038/nbt.4072.

Ching-Feng Chiu, Jonathan Jaime G Guerrero, Ric Ryan H Regalado, Ma Joy B Zamora, Jiayan Zhou, Kin Israel Notarte, Yu-Wei Lu, Paolo C Encarnacion, Cidne Danielle D Carles, Edrian M Octavo, et al. Insights into metabolic reprogramming in tumor evolution and therapy. *Cancers*, 16(20):3513, 2024.

Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35 (1):53–65, 2018.

Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.

Ali Ebrahim, Joshua A Lerman, Bernhard O Palsson, and Daniel R Hyduke. Cobrapy: Constraints-based reconstruction and analysis for python. *BMC Systems Biology*, 7(1), August 2013. ISSN 1752-0509. doi: 10.1186/1752-0509-7-74. URL http://dx.doi.org/10.1186/1752-0509-7-74.

Mohamed El-Tanani, Syed Arman Rabbani, Yahia El-Tanani, and Ismail I Matalka. Metabolic vulnerabilities in cancer: A new therapeutic strategy. *Critical Reviews in Oncology/Hematology*, pp. 104438, 2024.

Clémence Joseph, Haris Zafeiropoulos, Kristel Bernaerts, and Karoline Faust. Predicting microbial interactions with approaches based on flux balance analysis: an evaluation. *BMC bioinformatics*, 25(1):36, 2024.

Bin Liu, Bodo Rosenhahn, Thomas Illig, and David S DeLuca. A variational autoencoder trained with priors from canonical pathways increases the interpretability of transcriptome data. *PLOS Computational Biology*, 20(7):e1011198, 2024.

Kelly Osayi Otakhor and Elizabeth O Soladoye. A review of metabolic reprogramming in cancer cells: Mechanisms and therapeutic targets. *World Journal of Advanced Research and Reviews*, 23 (1):530–539, 2024.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.

Viktoria Schuster, Emma Dann, Anders Krogh, and Sarah A Teichmann. multidgd: A versatile deep generative model for multi-omics data. *Nature Communications*, 15(1):10031, 2024.

Aravind Subramanian, Pablo Tamayo, Vamsi K. Mootha, Sayan Mukherjee, Benjamin L. Ebert, Michael A. Gillette, Amanda Paulovich, Scott L. Pomeroy, Todd R. Golub, Eric S. Lander, and Jill P. Mesirov. Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences*, 102(43): 15545–15550, September 2005. ISSN 1091-6490. doi: 10.1073/pnas.0506580102. URL http://dx.doi.org/10.1073/pnas.0506580102.

Nektarios A Valous, Ferdinand Popp, Inka Zörnig, Dirk Jäger, and Pornpimol Charoentong. Graph machine learning for integrated multi-omics analysis. *British Journal of Cancer*, pp. 1–7, 2024.

Vishakha Vashisht, Ashutosh Vashisht, Ashis K Mondal, Jana Woodall, and Ravindra Kolhe. From genomic exploration to personalized treatment: Next-generation sequencing in oncology. *Current Issues in Molecular Biology*, 46(11):12527–12549, 2024.

Jiayuan Wang, QM Jonathan Wu, and Farhad Pourpanah. An attentive-based generative model for medical image synthesis. *International Journal of Machine Learning and Cybernetics*, 14(11): 3897–3910, 2023.

A More Analysis

A.1 Top Genes using SHAP

In this analysis provided in Figure 2 and 3, We first took the initial ten samples from our RNA-seq dataset to keep things straightforward and manageable. After computing the SHAP values for each sample's predicted metabolite fluxes, we averaged these values over the output dimension. This step gave us a single SHAP score for each gene, reflecting how much that gene, on average, influences the model's predictions across all flux outputs.

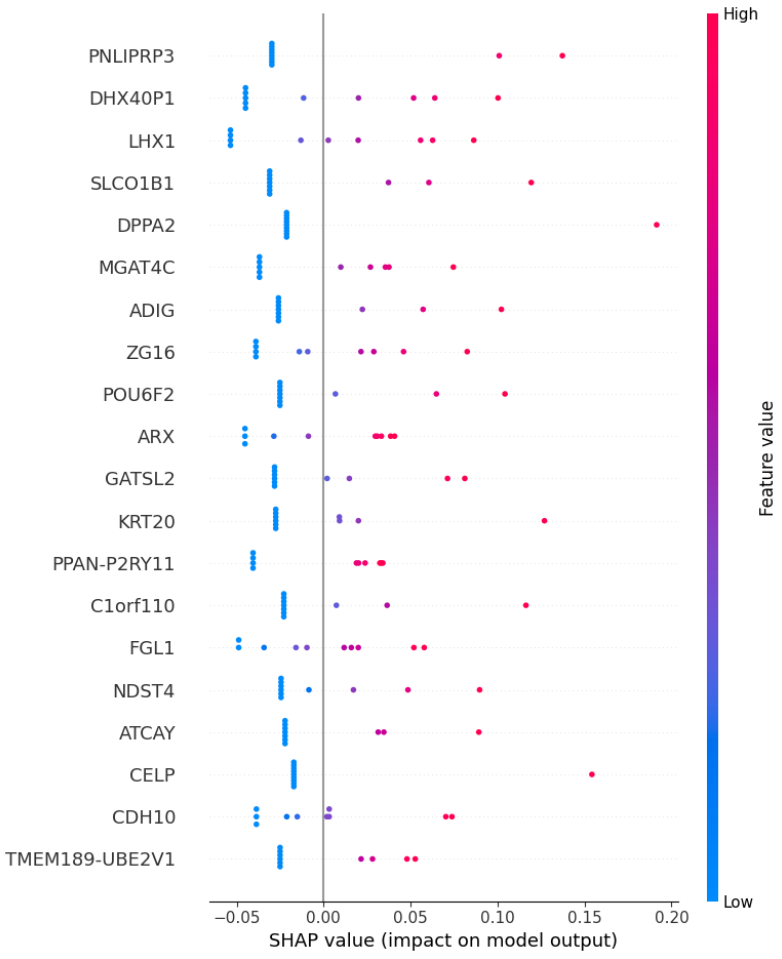

Figure 2: Top 10 Genes Contributing to Metabolite Prediction using SHAP Value Barplot

Next, we created a SHAP summary plot to visualize which genes are pushing the model's output higher or lower and to see how each gene's expression level (shown by color) affects the predicted flux. We then calculated the mean absolute SHAP values for every gene, picked the top ten with the highest scores, and displayed them in a bar chart. By focusing on these top contributors, we can

more easily pinpoint the genes that have the strongest overall impact on our metabolite predictions. This process helps us highlight key transcriptomic features for deeper biological investigation or for refining therapeutic strategies.

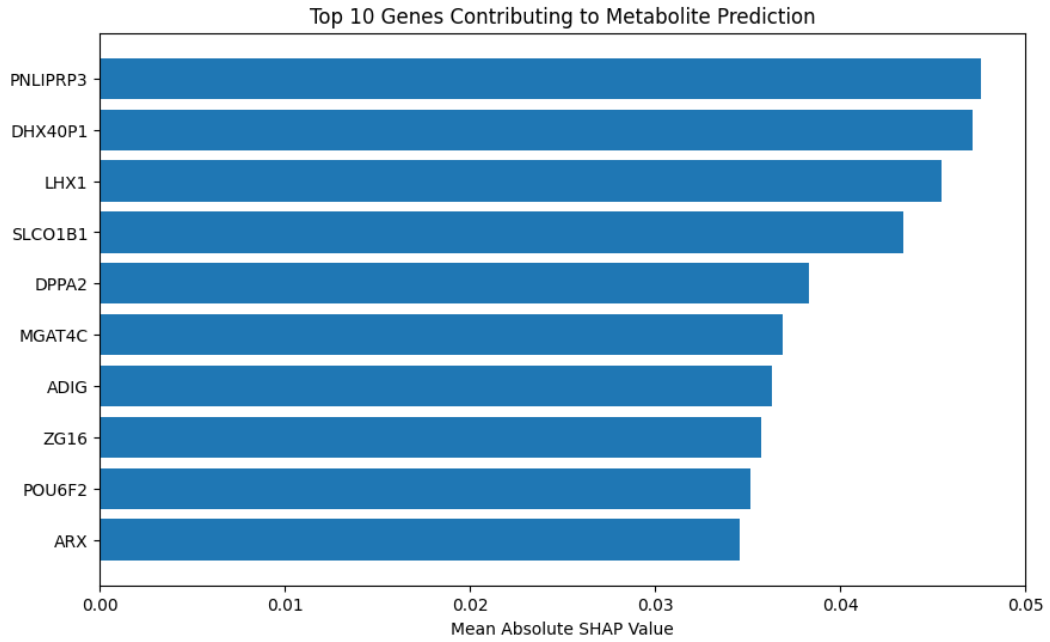

Figure 3: Top 10 Genes Contributing to Metabolite Prediction using Mean Absolute SHAP Value

A.2 SHAP Force Plot Analysis

We began by focusing on the first sample in our dataset and creating a shap Explanation object that captures how each gene influences the model's flux prediction. Specifically, we used the SHAP values for this sample, along with the expected (baseline) value and the sample's gene expression data, to visualize a force plot for each relevant output (in this case, Figure and). The force plot helps us see which genes are pushing the flux higher (shown in pink) or pulling it lower (shown in blue) compared to the baseline prediction. By comparing the two plots, we can see how different genes drive the flux in slightly different ways for each output. This view gives us a clearer picture of the local decision-making process of the model, allowing us to identify which transcriptomic features play the biggest roles in shaping the predicted metabolite fluxes for this particular sample.

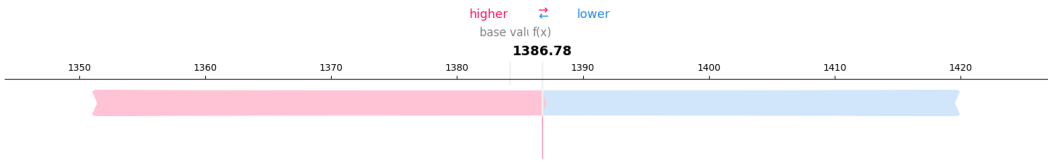

Figure 4: SHAP Force Plot for First RNA Sample

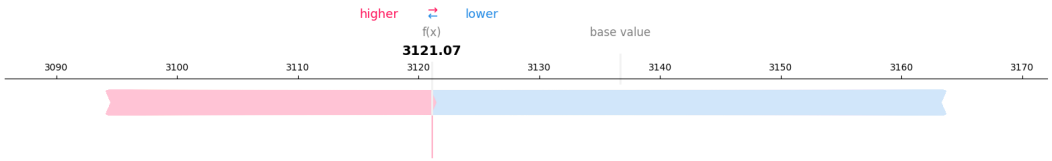

Figure 5: SHAP Force Plot for Second RNA Sample

A.3 SHAP Force Plot Analysis

To analyze the top enriched hallmark pathways, we use GSEA results, focusing on identifying the most significantly enriched pathways in our dataset. GSEA helps us understand which biological pathways, represented by sets of genes, are over-represented or under-represented in our data. By sorting the pathways based on their Combined Score, we can prioritize the pathways that are most significantly enriched. The bar chart in Figure reveals that the 'Pperoxisome' pathway has the highest Combined Score, indicating it's the most significantly enriched, suggesting a major role in metabolic functions like lipid metabolism and detoxification. Following this, 'Bile Acid Metabolism' and 'Apoptosis' are also highly enriched, pointing to changes in digestion, lipid processing, and cell death mechanisms. Other enriched pathways like 'Xenobiotic Metabolism', 'Adipogenesis', and 'Estrogen Response Late' highlight involvement in detoxification, fat cell formation, and estrogen signaling. This analysis helps us focus on key biological processes for further investigation.

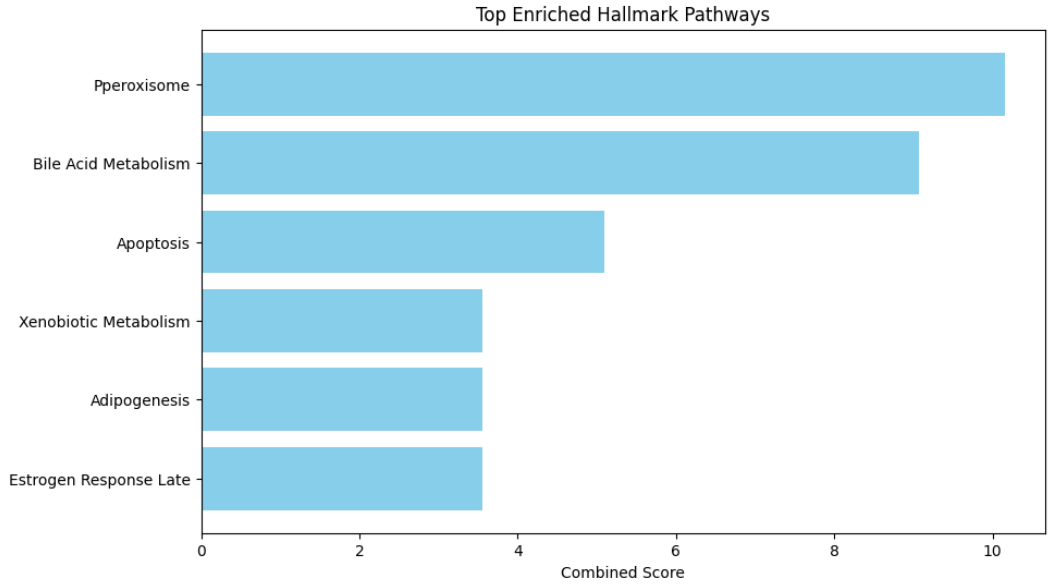

Figure 6: Top Enriched Hallmark Pathways

A.4 Analyzing the Impact of Radiation Therapy on Synthetic Glutathione Flux

In our study, we aimed to investigate how radiation therapy influences glutathione production, a crucial antioxidant in cells. To do this, we used our trained PathGAN model to predict synthetic metabolic fluxes from gene expression data. We fed the gene expression data into the trained gener-

ator to predict metabolic fluxes, extracting the glutathione flux values. We stored these values in a pandas DataFrame, ensuring each sample was correctly indexed.

Next, we merged the synthetic flux data with clinical information, particularly focusing on the radiation therapy status, to examine potential differences in glutathione flux between groups that did and did not receive radiation therapy. We then visualized these differences using a boxplot, with radiation therapy status on the x-axis and synthetic glutathione flux on the y-axis.

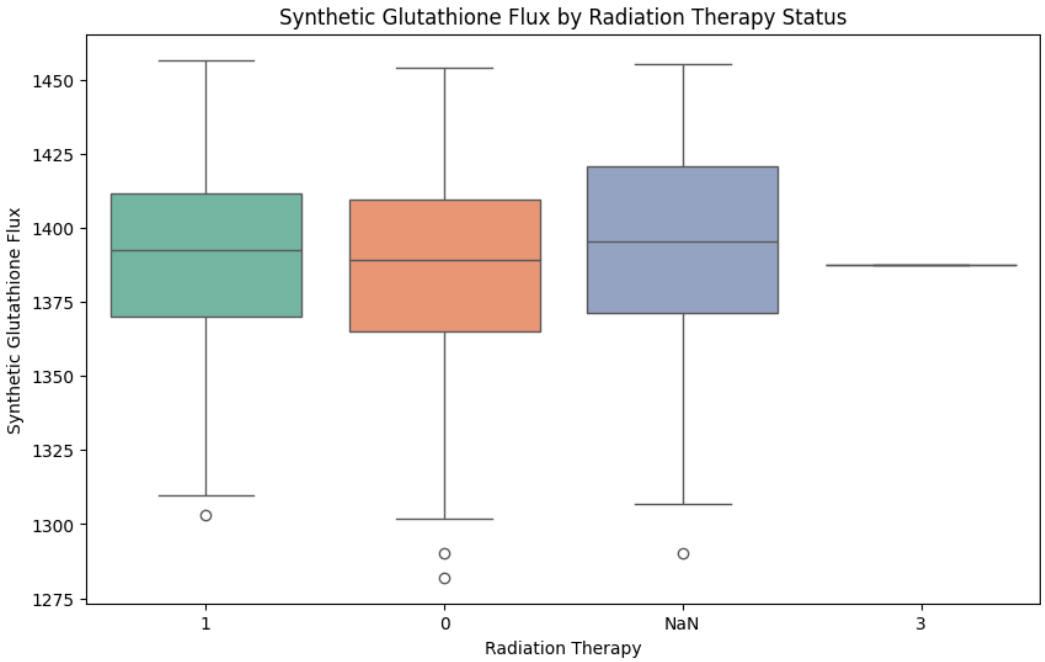

Figure 7: Synthetic Glutathione Flux by Radiation Therapy Status

The boxplot showed some differences in glutathione flux between the groups, particularly with the radiation therapy group ('1') displaying a higher median flux than the non-radiation therapy group ('0'). However, there was also overlap in distributions, suggesting that the differences might not be large. The 'NaN' and '3' groups showed more variability and potential outliers, indicating the need for further investigation into the biological factors affecting glutathione production. This visualization provides a starting point for understanding the relationship between radiation therapy and glutathione flux, with further statistical analysis needed for more definitive conclusions.