





Supplementary Figure S1. Calibration curve for the breast cancer dataset



Supplementary Figure S2. Label shuffling sanity check on the breast cancer dataset (100 runs)

Supplementary Table S1. Comparative analysis of fusion strategies (95% CI)

Dataset	Fusion strategy	Accuracy	Precision	Recall	ROC_AUC	PR_AUC
Breast cancer	ComplementaryFusion	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]
Breast cancer	ConcatFusion	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]
Breast cancer	AdaptiveGatedFusion	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]
Breast cancer	AttentionFusion	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]	1.000[1.000–1.000]
Gastric cancer	ComplementaryFusion	0.895[0.832–0.947]	0.908[0.833–0.970]	0.937[0.868–0.985]	0.956[0.909–0.990]	0.975[0.942–0.995]
Gastric cancer	ConcatFusion	0.851[0.779–0.916]	0.825[0.740–0.904]	0.983[0.945–1.000]	0.948[0.893–0.988]	0.969[0.928–0.994]
Gastric cancer	AdaptiveGatedFusion	0.872[0.800–0.937]	0.880[0.797–0.955]	0.935[0.868–0.985]	0.936[0.865–0.987]	0.955[0.893–0.994]
Gastric cancer	AttentionFusion	0.883[0.811–0.947]	0.870[0.789–0.944]	0.967[0.915–1.000]	0.953[0.895–0.992]	0.970[0.928–0.996]

ConcatFusion: This is the simplest approach, which directly concatenates the output vectors from the pathway subnetwork and the structure subnetwork. This method assumes that both branches contribute equally without any interaction or weighting.

AdaptiveGatedFusion: This strategy learns adaptive gating mechanisms at two levels. It computes element-wise gates to selectively strengthen or suppress individual feature dimensions in each branch, while also learning a global gate to control the overall contribution of the pathway branch versus the structure branch.

AttentionFusion: This method uses a lightweight neural network to learn sample-specific attention weights. It takes the combined features from both branches as input and dynamically assigns higher importance to the more relevant branch for each individual sample. The weighted features from both branches are then concatenated to form the final representation.

ComplementaryFusion: The proposed fusion module is specifically designed to capture the complementary relationship between metabolic pathway information and

molecular structural information. It first calculates a complementarity score that reflects how much the two branches can supplement each other for the current sample. This score is then used to enhance the representations from both branches. Additionally, an explicit complementary feature is generated to model their interactions. The enhanced features and the complementary feature are finally combined through a learnable mapping layer. This design encourages synergistic integration rather than simple competition or averaging between the two biological views.

Supplementary Table S2. Top-10 Node Overlap Rates per Layer (100 Bootstrap Resamples)

Branch	Layer	Average overlap rate	Minimum overlap rate	Maximum overlap rate	Number of nodes
Pathway	Layer0	0.89	0.70	1.00	132
Pathway	Layer1	0.96	0.90	1.00	208
Pathway	Layer2	1.00	1.00	1.00	90
Pathway	Layer3	1.00	1.00	1.00	26
Structure	Layer0	0.90	0.70	1.00	132
Structure	Layer1	0.93	0.90	1.00	51
Structure	Layer2	0.96	0.90	1.00	29
Structure	Layer3	1.00	1.00	1.00	8

Bootstrap resampling analysis (n=100) demonstrated that node importance rankings became progressively more stable in higher layers. At the top pathway layer (Layer 3, 26 nodes) and structural class layer (Layer 3, 8 nodes), Top-10 overlap rates reached 1.00 across all resampled subsets, indicating that the identification of the most important biological pathways and structural classes is entirely robust to sample variation.