

Bridging Morphology and Molecular Signatures: Multi-Task Deep Learning for Multi-Omics Prediction from Histopathology

Anonymous CVPR submission

Paper ID ****

001

1

Abstract

002 *Whole slide images (WSIs) capture intricate morpho-*
003 *logical features that correlate with molecular profiles in*
004 *tumors, making them valuable for non-invasive molecular*
005 *profiling. While previous work in computational pathology*
006 *has focused on predicting gene expression, protein levels,*
007 *and mutation status from WSIs, we extend this by introduc-*
008 *ing a deep learning framework that predicts pathway activ-*
009 *ity and microRNA (miRNA) expression in addition to these*
010 *commonly studied molecular layers. By employing multi-*
011 *task learning, our model efficiently captures shared histo-*
012 *logical patterns across different molecular modalities, en-*
013 *hancing prediction accuracy. We show that pathway activ-*
014 *ity is the most reliably predicted feature, followed by protein*
015 *expression, with gene expression and miRNA predictions*
016 *being more challenging. These findings highlight the neces-*
017 *sity of incorporating pathway and miRNA data for a more*
018 *comprehensive and biologically relevant understanding of*
019 *tumor biology. Our approach demonstrates significant po-*
020 *tential for improving cancer diagnostics and biomarker dis-*
021 *covery, offering a more comprehensive alternative to tradi-*
022 *tional molecular assays.*

023

1. Introduction

024 Hematoxylin and eosin (H&E)-stained histopathology
025 slides have long been fundamental in cancer diagnosis, of-
026 fering intricate insights into tissue architecture and cellu-
027 lar organization. Recent advances, however, have revealed
028 that these slides also contain rich molecular information that
029 correlates with alterations at the transcriptomic, proteomic,
030 pathway, and microRNA levels [4–6]. Harnessing this
031 latent information to predict omics profiles directly from
032 WSIs presents a promising, rapid, and cost-effective alter-
033 native to traditional molecular profiling methods, which are
034 often labor-intensive and expensive [8].

035 Deep learning techniques have revolutionized compu-
036 tational pathology, enabling the extraction of biologically

037 meaningful features from WSIs [14, 17, 21]. While
038 molecular profiling typically relies on specialized labora-
039 tory workflows, recent advancements in convolutional neu-
040 ral networks (CNNs), multiple-instance learning (MIL),
041 and transformer-based architectures have demonstrated that
042 histopathological patterns in WSIs can predict underlying
043 molecular phenotypes [2, 7, 10, 11, 19]. These techniques
044 have successfully been applied to predict gene expression,
045 protein abundance, and other molecular signatures from his-
046 tology, setting the stage for AI-driven virtual multi-omics
047 profiling [8].

Related works to predict omics level from WSI

048 Recent advances in deep learning have enabled the pre-
049 diction of molecular profiles (mainly gene expression, mu-
050 tation status and protein expression profiles directly from
051 WSIs offering a non-invasive alternative for molecular pro-
052 filing in histopathology. Existing approaches typically fol-
053 low a shared pipeline with variations in architectural com-
054 ponents. The standard workflow involves preprocessing
055 WSIs by segmenting and tiling them into square patches.
056 A subset of N patches is randomly sampled, and patch-
057 level embeddings are generated using a pre-trained feature
058 extractor. The WSI-level representation is then obtained by
059 aggregating these patch embeddings into a single latent vec-
060 tor $W \in R^{1 \times d}$ which is passed to a downstream predictor
061 often a shallow linear layer to map embeddings to molec-
062 ular profiles. For brevity, we refer to the combined tiling,
063 feature extraction, and aggregation modules as the encoder
064 part.

065 In WSI-based omics prediction, model architectures typ-
066 ically consist of two core components: feature extrac-
067 tion and patch aggregation. The feature extractor gener-
068 ates patch-level embeddings, while the aggregator com-
069 bines these embeddings into a WSI-level representation for
070 downstream omics prediction.
071

072 For feature extraction, recent approaches leverage both
073 CNN-based and transformer-based models, pre-trained on
074 large-scale histopathology datasets. Commonly used mod-
075 els include CTransPath[20], which uses a hierarchical
076 transformer architecture to capture local and global tissue

077	patterns, and UNI[3], a unified framework trained in multiple histology datasets for robust feature extraction. These foundation models effectively encode WSI patches into expressive feature representations, capturing morphological variations relevant to molecular profiles.	130
078		131
079		132
080		133
081		134
082	For patch aggregation, various strategies are employed to combine patch embeddings into a single WSI-level vector. MLP-based aggregators, as seen in HE2RNA[17], apply fully connected layers to directly merge patch features. Transformer-based aggregators, such as tRNAsformer[1], use self-attention to capture inter-patch relationships, enabling the model to account for spatial dependencies. More recent methods, like SEQUOIA[16] adopt SummaryMixing[15], a linearized transformer variant, to enhance efficiency while preserving representational power.	135
083		136
084		137
085		138
086		139
087		140
088		
089		
090		
091		
092	While prior research has primarily concentrated on predicting gene expression or mutation status from WSIs, with some extending to protein expression, the inference of pathway activity and microRNA (miRNA) expression has largely been unexplored. Pathway activity provides a more interpretable representation of molecular states compared to individual gene expression, capturing the collective behavior of genes involved in biological processes. This makes pathway-based inference not only more robust but also more biologically meaningful for downstream analysis. Similarly, miRNAs, which regulate gene expression post-transcriptionally, are vital biomarkers for cancer detection and prognosis, yet their prediction from WSIs remains largely unexplored. Moreover, most existing models focus on predicting individual molecular modalities, such as gene expression or proteomics, independently. However, the molecular processes in cells are intricately interconnected, with gene expression, protein regulation, and pathway activity influencing one another. By ignoring these interdependencies, current models may fail to capture biologically meaningful cross-omics relationships, thereby limiting their overall predictive accuracy.	141
093		
094		
095		
096		
097		
098		
099		
100		
101		
102		
103		
104		
105		
106		
107		
108		
109		
110		
111		
112		
113		
114	In this work, we expand the scope of WSI-based molecular inference to include multiple omics layers, including gene expression, protein expression, pathway activity, and miRNA levels. Through both single-task and multi-task learning frameworks, we demonstrate that pathway and miRNA predictions not only enhance interpretability but also provide valuable prognostic insights. Our findings suggest that inferring higher-level molecular features, such as pathway activity, offers greater robustness and clinical relevance, thus advancing the integration of histology and molecular data in cancer research.	142
115		143
116		144
117		145
118		146
119		147
120		148
121		149
122		150
123		151
124		152
125	By incorporating <i>pathway</i> and <i>microRNA</i> expression levels into WSI-based omics prediction, the framework moves beyond gene-level inference, providing a more stable, interpretable, and biologically relevant representation of tumor biology. This expanded approach enhances the utility	153
126		154
127		155
128		156
129		157
		158
		159
		160
		161
		162
		163
		164
		165
		166
		167
		168
		169
		170
		171
		172
		173
		174
		175
		176
		177
		178
		179

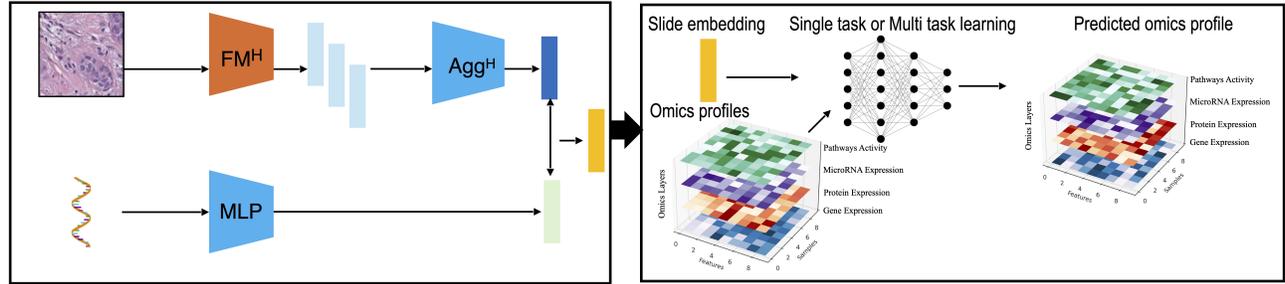


Figure 1. A histology slide is divided into patches and encoded using a pre-trained UNI encoder (FM^H). The patch embeddings are aggregated by an ABMIL module (Agg^H) into a slide-level representation. In parallel, gene expression data is encoded with an MLP. A SymCLR objective aligns the embeddings. The learned slide embeddings are then used to train deep regression models in single-task or multi-task settings to predict multi-omics profiles, including gene, protein, pathway, and microRNA expression.

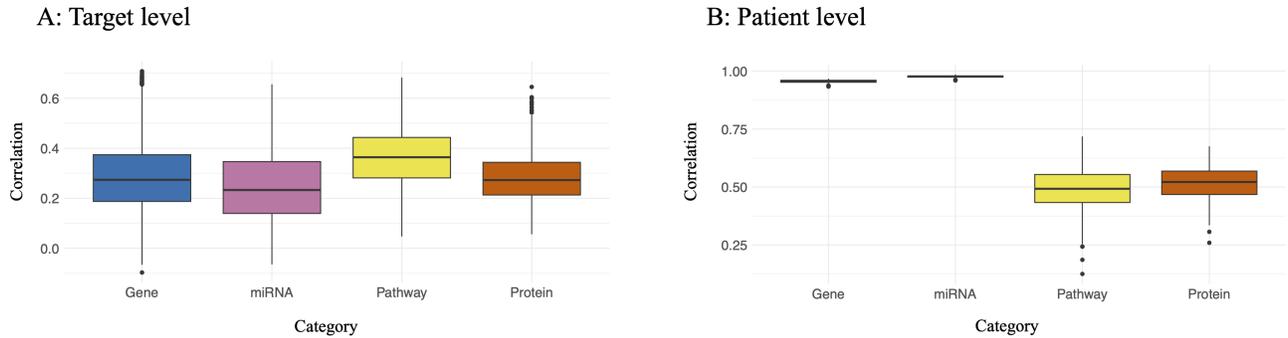


Figure 2. Model performance comparison across omics layers. (A) Correlation between predicted and actual omics profiles across patients. (B) Correlation of model performance at the individual patient level.

$$f_g : R^d \mapsto R^{m_g}$$

where m_g denotes the number of genes. Similarly, for pathway activity and protein expression prediction, we define:

$$f_p : R^d \mapsto R^{m_p}, f_k : R^d \mapsto R^{m_k}, f_l : R^d \mapsto R^{m_l}$$

180 where m_p and m_k represent the number of pathways and
 181 proteins, respectively. Each function f is parameterized as
 182 a fully connected neural network with a single hidden layer
 183 of 512 neurons, activated by ReLU, followed by an output
 184 layer tailored to the corresponding omics dimension. The
 185 model is trained independently for each omics type using
 186 the *Adam* optimizer and a mean squared error (MSE) loss
 187 function, ensuring accurate reconstruction of omics pro-
 188 files from histopathological features. Model selection is
 189 performed using 5-fold cross-validation, where the model
 190 achieving the lowest validation MSE is selected for final
 191 evaluation.

192 **Multi-task learning** Building on our single-task frame-
 193 work, we extend our approach to multi-task learning (MTL)

194 to jointly predict gene expression, pathway activity, and
 195 protein expression from WSI embeddings. Rather than
 196 training independent models for each molecular layer, we
 197 introduce a multi-task architecture that learns shared rep-
 198 resentations while allowing task-specific adaptations. This
 199 enables the model to leverage common histopathological
 200 features while maintaining flexibility to capture distinct pat-
 201 terns relevant to each omics type. Given a WSI embedding
 202 $z \in R^d$ extracted using TANGLE (Fig.1), we define a joint
 203 function:

$$f_\theta : R^d \mapsto (R^{m_g}, R^{m_p}, R^{m_k}, R^{m_m})$$

204 where m_g, m_p, m_k, m_m correspond to the number of
 205 genes, pathways, proteins, and microRNAs respectively.

206 Each omics prediction task is modeled as a high-
 207 dimensional regression problem, where a shared feature ex-
 208 traction backbone is followed by task specific output lay-
 209 ers. The shared backbone consists of a fully connected (FC)
 210 layer with 512 neurons and ReLU activation, capturing a
 211 common histological representation of tumor morphology.
 212 The task-specific heads, one for each omics modality, con-
 213 sist of independent linear layers, each mapping the features

	Top 50		Top 500	
	Single Task	Multi Task	Single Task	Multi Task
Gene	0.6778 \pm 0.0226	0.6741 \pm 0.0281	0.5573 \pm 0.0244	0.5935 \pm 0.0194
Protein	0.4694 \pm 0.0218	0.4924 \pm 0.0088	0.2897 \pm 0.0161	0.2958 \pm 0.0127
microRNA	0.5295 \pm 0.0278	0.5338 \pm 0.0170	0.2624 \pm 0.0243	0.3259 \pm 0.0099
Pathways	0.5970 \pm 0.0071	0.6449 \pm 0.0263	0.4902 \pm 0.0207	0.5031 \pm 0.0187

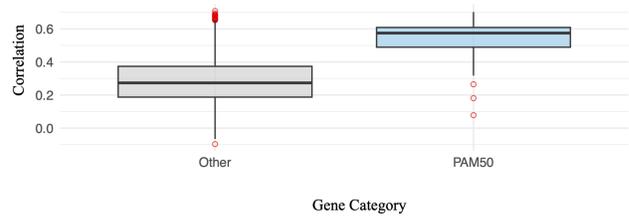


Figure 3. Comparison of single task vs. multi task learning and the prediction performance of breast cancer-relevant genes (PAM50). The left table presents quantitative results (mean Pearson correlation \pm standard deviation) for the top 50 and top 500 molecular targets across different omics modalities, while the right panel visualizes model performance on PAM50 gene signatures.

214 to the corresponding output dimension.

$$f_g(z) = W_g \cdot z + b_g, \quad f_p(z) = W_p \cdot z + b_p,$$

$$f_k(z) = W_k \cdot z + b_k, \quad f_m(z) = W_m \cdot z + b_m$$

where W and b are task-specific weight matrices and biases.

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \sum_{t \in \{g,p,k,m\}} M_t^i \cdot \|y_t^i - f_t(z^i)\|^2$$

215 To address missing omics measurements across patients,
 216 we employ task-specific masking, where a binary mask is
 217 applied during loss computation to exclude missing tar-
 218 gets. The model is trained using a partial supervision strat-
 219 egy, ensuring that only available omics labels contribute
 220 to the gradient update. Optimization is performed using
 221 the Adam optimizer, with the loss function defined as the
 222 masked mean squared error (MSSE) across all available
 223 omics tasks. Model selection is conducted via 5-fold cross-
 224 validation, selecting the model with the lowest validation
 225 MSE for final evaluation.

226 3. Experiments and results

227 3.1. Dataset

228 This study utilizes publicly available multimodal data from
 229 the Xena browser for breast cancer (BRCA)[9]. We col-
 230 lected a dataset comprising matched gene expression pro-
 231 files, WSIs, protein expression levels, microRNA expres-
 232 sion, and pathway activity measurements for a cohort of
 233 1,024 BRCA patients.

234 **Histology slides:** Matched histopathology slides for
 235 BRCA were obtained from the Xena browser and processed
 236 using the TANGLE framework. This approach generated
 237 patch embeddings, enabling the model to learn meaningful
 238 representations of tissue morphology for downstream pre-
 239 dictive tasks.

240 **Gene expression data:** We obtained matched gene ex-
 241 pression data with FPKM values for protein-coding genes

from the Xena browser. To improve the interpretability of
 the model, we excluded genes with a median expression of
 zero, retaining 17,759 protein-coding genes. To mitigate
 the dominance of highly expressed genes in the regression
 analysis, we applied a $\log_{10}(1+a)$ transformation to ex-
 pression values.

microRNA expression data: We obtained matched mi-
 croRNA expression data with FPKM values for 1800 mi-
 croRNA from the Xena browser. To enhance the inter-
 pretability of the model, we excluded microRNA with a me-
 dian expression of zero, retaining 751 microRNA genes. To
 mitigate the dominance of highly expressed genes in the re-
 gression analysis, we applied a $\log_{10}(1+a)$ transformation
 to expression values.

Protein expression data: We obtained proteomics
 data for breast cancer across 980 samples from the
 Xena Browser, measured using reverse-phase protein array
 (RPA), which includes expression levels for 480 proteins.
 Missing values were imputed using the median expression
 of the respective protein across all samples, ensuring data
 completeness. To stabilize variance and improve down-
 stream analysis, the entire dataset was subsequently loga-
 rithmically transformed.

Pathway activity data: We obtained pathway activity
 data for breast cancer from the Xena Browser, using the
 PARADIGM algorithm for inference[18]. PARADIGM
 integrates pathway information with gene expression and
 copy number variation data to estimate the activation status
 of pathway components within a unified pathway network.
 This network structure consolidates 1,387 constituent path-
 ways derived from three major pathway databases: NCI-
 PID, BioCarta, and Reactome, providing a comprehensive
 framework for pathway activity analysis. Furthermore, we
 evaluated TANGLE in both multimodal mode (contrastive
 alignment) and unimodal mode. However, we did not ob-
 serve a significant improvement in omics layer prediction
 between the two settings, suggesting that multimodal align-
 ment did not substantially enhance predictive performance
 in this context.

281 3.2. Single task learning of multi-omics

282 We performed omics prediction at both the target and pa-
283 tient levels and evaluated the model’s performance across
284 different molecular layers. At the target level, we assessed
285 the accuracy of predicting individual gene, protein, path-
286 way, and microRNA expression levels, while at the patient
287 level, we predicted the entire omics profile for each patient.
288 The predictive performance of our model varied between
289 different omics layers inferred from WSI of breast cancer
290 tissues Figure 2. Pathway activity exhibited the highest
291 correlation between predicted and hold-out data, followed
292 by protein expression, while miRNA and gene expression
293 showed comparatively lower correlations (Figure 2).

294 **Pathway activity prediction:** The superior performance
295 in pathway prediction is likely due to the aggregation of
296 gene expression into functional networks. By capturing col-
297 lective biological activity rather than individual gene fluc-
298 tuations, this approach reduces transcriptional noise and
299 stochastic variability, leading to more stable and inter-
300 pretable predictions. The top predicted pathways that were
301 strongly linked to cancer progression and relevant biolog-
302 ical processes include Targets of *MYC* transcriptional acti-
303 vation, *Aurora B signaling*, and *Cyclin B2* mediated events.
304 These pathways play pivotal roles in cell cycle regulation,
305 tumor invasion, and cellular growth. *C-MYC*, a key onco-
306 gene, activates several transcriptional programs that drive
307 tumor genesis and progression, including genes involved
308 in cellular metabolism, proliferation, and survival. The
309 model’s ability to predict pathway activity associated with
310 *C-MYC* transcriptional activation implies that the morpho-
311 logical features of WSI reflect key signaling events that
312 drive the pathology of breast cancer. *Aurora B signaling* is
313 involved in the regulation of mitosis and chromosome seg-
314 regation, crucial for proper cell division and genomic sta-
315 bility. These pathway predictions suggest that WSI cap-
316 tures cellular processes that underlie tumor proliferation
317 and metastasis, providing clinically relevant insights into
318 cancer progression at the tissue level.

319 **Protein expression prediction:** showed stronger corre-
320 lation compared to microRNA and gene expression. Protein
321 expression prediction from WSI is inherently more accurate
322 than gene expression inference due to the stronger morpho-
323 logical imprint of proteins on tissue architecture. Unlike
324 mRNA, which reflects transient transcriptional states, pro-
325 tein abundance integrates post-transcriptional and transla-
326 tional regulation, resulting in more stable and visually dis-
327 cernible patterns. Proteins directly influence cellular mor-
328 phology, adhesion, and staining intensity, creating distinct
329 histological signatures that deep learning models can ef-
330 fectively capture. In contrast, gene expression exhibits
331 higher variability due to transcriptional noise, cell-type het-
332 erogeneity, and dynamic regulation, making it less reliably
333 linked to tissue morphology. Additionally, the lower dimen-

sional- 334
ity of proteomics data compared to transcriptomics 335
reduces prediction complexity, further improving model ac- 336
curacy. These factors collectively make protein expression 337
more amenable to robust inference from WSI in computa- 338
tional pathology.

339 **miRNA Prediction:** despite its moderate accuracy, the 339
model’s ability to predict miRNA expression from WSI re- 340
veals that breast cancer tissue morphology encodes subtle 341
yet detectable imprints of miRNA activity. miRNAs are 342
critical regulators of post-transcriptional gene expression, 343
influencing key cancer processes such as cell proliferation, 344
apoptosis, and metastasis. Their successful inference from 345
histological data suggests that the morphological features 346
of tumor tissues capture molecular dysregulation associated 347
with miRNA activity, providing a potential new biomarker 348
for histology-based cancer profiling. However, the rela- 349
tively lower accuracy of miRNA prediction compared to 350
protein and pathway activity can be attributed to the com- 351
plexity of miRNA regulation, which is highly tissue-specific 352
and often present at lower abundance. These factors make 353
it more challenging to directly capture miRNA expression 354
from morphological patterns alone. Nonetheless, the abil- 355
ity to predict miRNA expression from WSI opens up ex- 356
citing possibilities for incorporating small RNA data into 357
histopathological analyses, enhancing the predictive capa- 358
bilities of cancer diagnostics and enabling more compre- 359
hensive molecular profiling from tissue slides. 360

361 **Gene expression prediction:** Gene expression predic- 361
tion from WSI showed the lowest correlation among the 362
evaluated omics layers. This lower performance can be 363
attributed to several challenges, including technical noise, 364
cell-type heterogeneity, and regulatory mechanisms that ex- 365
tend beyond the morphological features captured in WSI. 366
Since gene expression is influenced by complex, tissue- 367
specific regulatory networks, predicting expression for all 368
genes from histological data alone is inherently more dif- 369
ficult. The morphological features of tissue sections pro- 370
vide only indirect insights into the transcriptomic land- 371
scape, leading to reduced accuracy compared to protein 372
and pathway predictions, where functional relationships are 373
more directly reflected in tissue morphology. This high- 374
lights the limitations of using WSI for gene expression- 375
based biomarker inference, which may be suboptimal and 376
a limiting factor. This underscores the need to model other 377
omics layers, such as pathways, proteins, or small RNA, 378
to infer more accurate and comprehensive predictions from 379
WSI. 380

381 **Patient level prediction:** beyond target-level predic- 381
tions, our model demonstrates superior performance in ac- 382
curately predicting omics profiles at the patient level. This 383
can likely be attributed to the inherent correlations within 384
each omics modality—genes within the transcriptomics 385
layer are correlated with each other, and similarly, pro- 386

teins within the proteomics layer exhibit strong correlations. When the model performs at the patient level, it can leverage these correlation structures, which helps improve predictive accuracy. In contrast, at the target level, the model struggles to utilize these correlations since the focus is on individual targets like specific genes or proteins, rather than on the collective relationships between them. This lack of structural correlation at the target level limits the model's ability to make accurate predictions across modalities.

3.3. Multi task learning of multi-omics

Using the architecture described in the Methods section, we conducted multi-task learning (MTL) across multiple settings. Initially, we trained a unified model to predict 17759 genes, 778 microRNAs, 1857 pathways, and 457 proteins, following the STL baseline. However, we observed no significant performance improvements with MTL. A key factor could be the imbalance in task dimensionality, where large output spaces dominate optimization, leading to suboptimal representation learning for smaller tasks. To mitigate this, we constrained MTL to the top 500 genes, pathways, microRNAs, and proteins, aligning task dimensions and reducing the risk of unequal loss scaling. The 480-protein constraint naturally set 500 as a threshold, harmonizing the learning dynamics across tasks.

To further refine task balance, we trained MTL on the top 50 genes, pathways, and proteins. This was motivated by the hypothesis that reducing dimensionality while retaining biologically relevant targets would enhance feature sharing and optimization stability. From an AI/ML perspective, high-dimensional MTL can suffer from negative transfer, where conflicting gradients degrade task-specific learning. By restricting MTL to a compact, high-signal subset, we aimed to improve cross-task knowledge transfer, stabilize shared representations, and prevent task interference.

Table 3 demonstrates the effectiveness of multi-task learning (MTL) across different omics modalities. While MTL and single-task learning (STL) showed similar performance for the top 50 features, MTL consistently outperformed STL in the top 500 setting, particularly for genes, microRNAs, and pathways. Notably, MTL improved gene prediction by 6.5% (0.5573 \rightarrow 0.5935) and microRNA prediction by 24% (0.2624 \rightarrow 0.3259), highlighting its ability to leverage shared representations across omics layers. Pathway prediction also benefited from MTL, showing an increase from 0.4902 to 0.5031, indicating that incorporating cross-task dependencies enhances predictive power. While protein prediction showed only a marginal improvement (0.2897 \rightarrow 0.2958), this suggests that further architectural refinements may be needed to better capture protein interactions in the MTL framework. Overall, these results demonstrate that MTL effectively leverages shared histological representations to improve multi-omics predic-

tion when task dimensionalities are aligned. The significant performance gains in gene and microRNA prediction suggest that incorporating structured multi-task objectives facilitates cross-task knowledge transfer, particularly for omics layers with overlapping regulatory mechanisms. These findings reinforce the potential of MTL for enhancing histology-based biomarker discovery by integrating complementary molecular signals.

3.4. Survival analysis

To evaluate the clinical relevance of the inferred omics profiles, we examined whether the top predicted targets were associated with patient survival outcomes. Specifically, we performed survival analysis on AGR, FOXM1, and HSA-miR-190b expression levels in breast cancer (BRCA) patients. The expression of these biomarkers was accurately predicted from WSI of BRCA tumors using our deep learning model (Fig. 4). Kaplan-Meier survival curves revealed significant associations between the predicted expression levels and patient survival.

AGR (Activator of G-protein signaling) was strongly linked to poorer overall survival (OS), suggesting its potential role in driving BRCA progression. Similarly, FOXM1 (Forkhead Box M1), a key regulator of cell cycle and proliferation, was associated with reduced survival, highlighting its value as a prognostic marker. In contrast, HSA-miR-190b, a tumor-suppressive microRNA, showed a protective effect, with higher expression correlating with improved survival outcomes. These results demonstrate that WSI-derived omics predictions can capture clinically relevant biomarkers, offering potential for non-invasive risk stratification and personalized treatment planning.

4. Discussion

In this study, we demonstrated the ability to predict multiple omics layers directly from WSI of breast cancer tissues using a deep learning framework. The predictive performance of our model varied across the different molecular modalities, with pathway activity emerging as the most predictable feature, followed by protein expression. miRNA and gene expression predictions showed lower correlations, highlighting the challenges and opportunities in histology-driven molecular inference.

The superior performance of pathway-level predictions suggests that aggregating gene expression into functional networks enhances predictive robustness. Biological pathways, by their nature, integrate signals from multiple genes that contribute collectively to a functional process. This aggregation helps to reduce transcriptional noise and stochastic variability, which are often associated with individual gene expression measurements. As a result, pathway-level inference offers more stable and interpretable predictions,

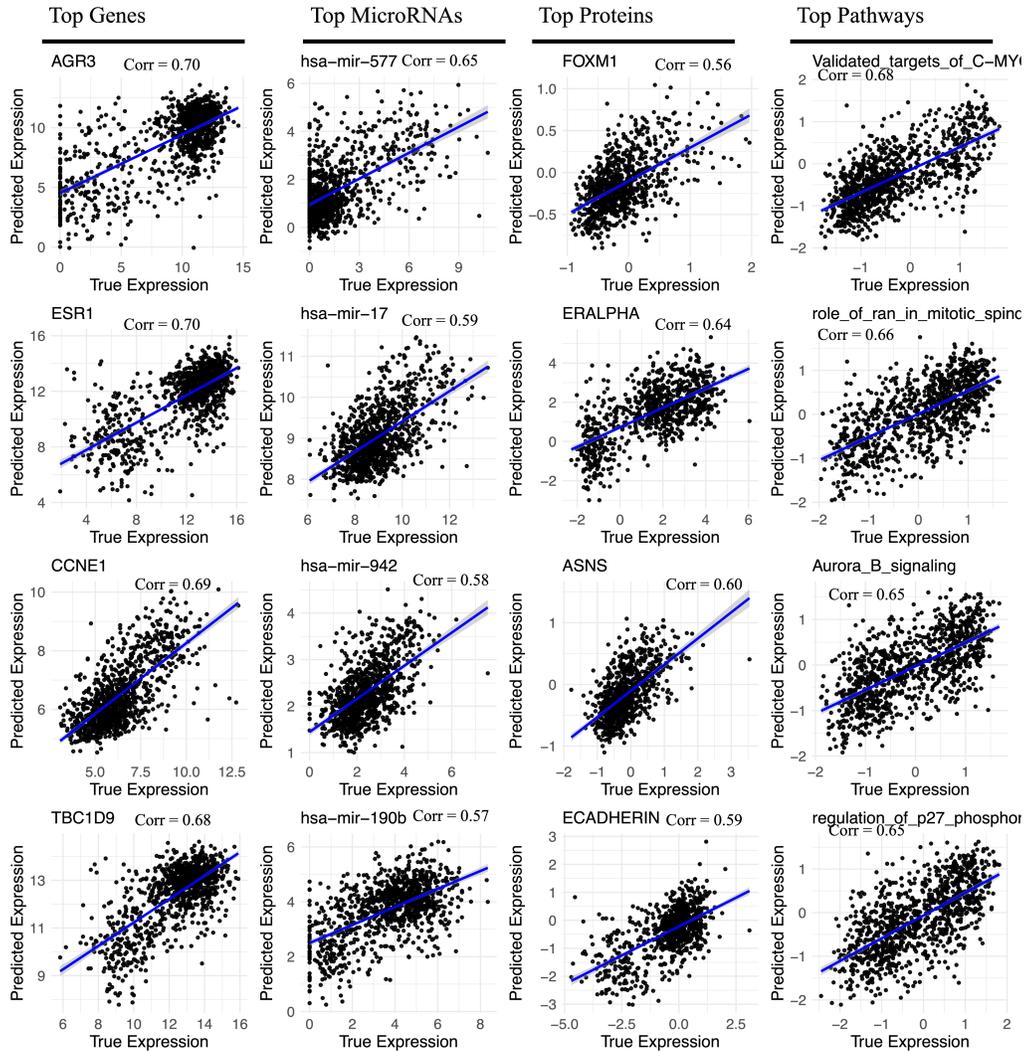


Figure 4. Correlation plot showing predicted vs. ground truth expression or activity values for the top 4 genes, proteins, microRNAs, and pathways.

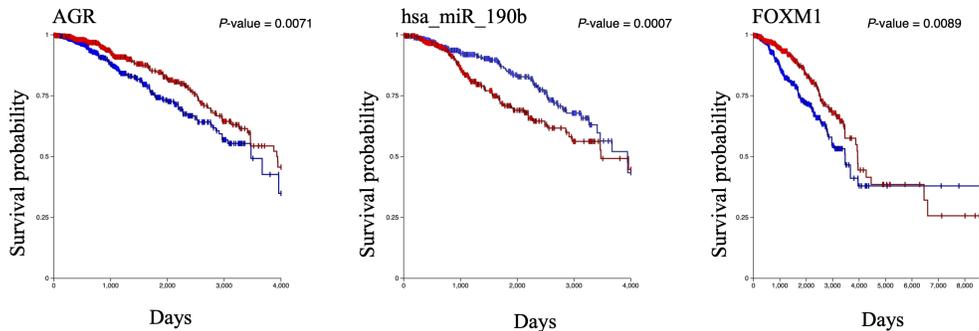


Figure 5. Survival analysis of AGR, FOXM1, and HSA-miR-190b expression levels in breast cancer (BRCA) patients, highlighting their prognostic significance.

489	effectively mitigating the challenges of capturing transcrip-	542
490	tional variability.	543
491	Protein expression prediction from WSI also demon-	544
492	strated a higher correlation compared to miRNA and gene	545
493	expression. This finding indicates that histological features,	546
494	which reflect tissue morphology, are better at capturing as-	547
495	pects of post-transcriptional regulation and protein abun-	548
496	dance than transcriptomic data alone. Since protein ex-	549
497	pression levels are more directly influenced by cellular ar-	550
498	chitecture and histological features, it is not surprising that	551
499	these predictions were more accurate than gene expression	552
500	predictions. The intermediate performance of miRNA pre-	553
501	diction underscores the potential for leveraging histolog-	554
502	ical patterns to infer regulatory elements beyond mRNA	555
503	and protein levels. While the prediction of miRNA ex-	556
504	pression from WSI is still a relatively novel approach, the	557
505	model's moderate success in this area suggests that miRNA-	558
506	associated regulatory activity leaves a morphological im-	559
507	print in cancer tissues, offering a promising avenue for fu-	560
508	ture exploration.	561
509	In contrast, gene expression prediction exhibited the	562
510	lowest performance in our study. This may be attributed to	563
511	several factors, including the greater variability of transcript	564
512	levels due to technical noise, cell-type heterogeneity, and	565
513	complex regulatory mechanisms that are not fully captured	
514	by histological features alone. The relatively lower corre-	566
515	lation for miRNA prediction compared to protein and path-	
516	way predictions further suggests that miRNA expression,	567
517	despite its biological relevance, is more difficult to infer	568
518	from histology. miRNA regulation is highly tissue-specific	569
519	and influenced by a complex network of post-transcriptional	570
520	processes, making it more challenging to detect directly	571
521	from morphological patterns. Nevertheless, the ability to	572
522	predict miRNA expression from WSI represents a signif-	573
523	icant advancement in integrating small RNA profiles into	574
524	histopathological analysis, contributing to a more compre-	575
525	hensive molecular understanding of cancer biology.	576
526	Beyond predictive accuracy, our study explored the bi-	
527	ological relevance of the model's predictions by focusing	577
528	on the PAM50 gene signature, a widely used panel for	
529	molecular subtyping in breast cancer. Our findings re-	578
530	veal that the model achieved significantly higher predic-	579
531	tion accuracy for the PAM50 genes compared to the rest	580
532	of the transcriptome ³ . This result underscores the model's	
533	ability to capture clinically relevant signals from histology	581
534	slides, as the PAM50 genes are tightly associated with spe-	
535	cific breast cancer subtypes that exhibit distinct histopatho-	582
536	logical patterns. The high prediction accuracy for these	583
537	genes highlights the potential of using WSI-based mod-	584
538	els for molecular subtyping, which could enhance clinical	585
539	decision-making in personalized treatment strategies.	586
540	We conducted survival analysis on the top predicted	587
541	genes, proteins, and miRNAs (AGR, FOXM1, and	588
	hsa-miR-190b) to evaluate their prognostic significance.	589
	Kaplan-Meier survival curves revealed strong associations	590
	between their expression levels and patient outcomes.	
	AGR, a regulator of G-protein signaling, correlated with	
	poorer overall survival, suggesting its role in cancer pro-	
	gression. Similarly, FOXM1, a key driver of cell cycle reg-	
	ulation and proliferation, was linked to reduced survival, re-	
	inforcing its potential as a prognostic marker. Conversely,	
	lower expression of hsa-miR-190b, a tumor-suppressive mi-	
	croRNA, was associated with worse survival outcomes, un-	
	derscoring its protective role. These findings highlight	
	the clinical relevance of histology-derived molecular pre-	
	dictions and their potential for risk stratification and per-	
	sonalized treatment strategies in breast cancer. Our study	
	highlights the advantage of pathway-based inference and	
	multi-omics integration in improving predictive accuracy	
	and clinical relevance. Pathway-level predictions offer a	
	more stable and interpretable representation of molecular	
	states, while protein and miRNA predictions provide valu-	
	able insights into post-transcriptional regulation. The abil-	
	ity to predict these features directly from WSI opens new	
	possibilities for integrating histological data with molecu-	
	lar profiling, offering a more accessible and cost-effective	
	alternative to traditional molecular assays.	
	5. Conclusion	
	This work extends molecular inference from WSIs beyond	
	gene and protein expression to include pathway activity and	
	miRNA expression, providing a more holistic view of tumor	
	biology. By leveraging MTL, we demonstrate that multi-	
	omics integration enhances prediction accuracy, particu-	
	larly for pathway activity and microRNAs, which are criti-	
	cal in cancer progression. These findings highlight the po-	
	tential for deep learning frameworks to transform histology-	
	based molecular profiling into a non-invasive, scalable ap-	
	proach for biomarker discovery and precision oncology.	
	6. Code Availability	
	The implementation of our framework, including data pre-	
	processing and model training scripts, is available upon re-	
	quest.	
	References	
	[1] Areej Alsaafin, Amir Safarpour, Milad Sikaroudi, Jason D	
	Hipp, and HR Tizhoosh. Learning to predict rna sequence	
	expressions from whole slide images with applications for	
	search and classification. <i>Communications Biology</i> , 6(1):	
	304, 2023. 2	
	[2] Salim Arslan, Julian Schmidt, Cher Bass, Debapriya Mehro-	
	tra, Andre Geraldes, Shikha Singhal, Julius Hense, Xiusi	
	Li, Pandu Raharja-Liu, Oscar Maiques, et al. A systematic	
	pan-cancer study on deep learning-based prediction of multi-	

- 591 omic biomarkers from routine pathology images. *Communi-* 648
592 *cations Medicine*, 4(1):48, 2024. 1 649
- 593 [3] Richard J Chen, Tong Ding, Ming Y Lu, Drew FK 650
594 Williamson, Guillaume Jaume, Andrew H Song, Bowen 651
595 Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, 652
596 et al. Towards a general-purpose foundation model for 653
597 computational pathology. *Nature Medicine*, 30(3):850–862, 654
598 2024. 2 655
- 599 [4] Xingjian Chen, Jiecong Lin, Yuchen Wang, Weitong Zhang, 656
600 Weidun Xie, Zetian Zheng, and Ka-Chun Wong. He2gene: 657
601 image-to-rna translation via multi-task learning for spatial 658
602 transcriptomics data. *Bioinformatics*, 40(6):btac343, 2024. 659
603 1 660
- 604 [5] Nicolas Coudray, Paolo Santiago Ocampo, Theodore Sakel- 661
605 laropoulos, Navneet Narula, Matija Snuderl, David Fenyő, 662
606 Andre L Moreira, Narges Razavian, and Aristotelis Tsirigos. 663
607 Classification and mutation prediction from non–small cell 664
608 lung cancer histopathology images using deep learning. *Nature* 665
609 *medicine*, 24(10):1559–1567, 2018. 666
- 610 [6] Muhammad Dawood, Mark Eastwood, Mostafa Jahanifar, 667
611 Lawrence Young, Asa Ben-Hur, Kim Branson, Louise Jones, 668
612 Nasir Rajpoot, et al. Cross-linking breast tumor transcrip- 669
613 tomic states and tissue histology. *Cell Reports Medicine*, 4 670
614 (12), 2023. 1 671
- 615 [7] Andreas Ekholm, Yinxi Wang, Johan Vallon-Christersson, 672
616 Constance Boissin, and Mattias Rantalainen. Prediction 673
617 of gene expression-based breast cancer proliferation scores 674
618 from histopathology whole slide images using deep learning. 675
619 *BMC cancer*, 24(1):1–13, 2024. 1 676
- 620 [8] Omar SM El Nahhas, Chiara ML Loeffler, Zunamys I Car- 677
621 rero, Marko van Treeck, Fiona R Kolbinger, Katherine J He- 678
622 witt, Hannah S Muti, Mara Graziani, Qinghe Zeng, Julien 679
623 Calderaro, et al. Regression-based deep-learning predicts 680
624 molecular biomarkers from pathology slides. *nature com-* 681
625 *munications*, 15(1):1253, 2024. 1 682
- 626 [9] Mary J Goldman, Brian Craft, Mim Hastie, Kristupas 683
627 Repečka, Fran McDade, Akhil Kamath, Ayan Banerjee, Yun- 684
628 hai Luo, Dave Rogers, Angela N Brooks, et al. Visualizing 685
629 and interpreting cancer genomics data via the xena platform. 686
630 *Nature biotechnology*, 38(6):675–678, 2020. 4 687
- 631 [10] Max Hallemeesch, Marija Pizurica, Paloma Rabaey, Olivier 688
632 Gevaert, Thomas Demeester, and Kathleen Marchal. Prior 689
633 knowledge injection into deep learning models predicting 690
634 gene expression from whole slide images. *arXiv preprint* 691
635 *arXiv:2501.14056*, 2025. 1
- 636 [11] Danh-Tai Hoang, Gal Dinstag, Eldad D Shulman, Leandro C 692
637 Hermida, Doreen S Ben-Zvi, Efrat Elis, Katherine Caley, 693
638 Stephen-John Sammut, Sanju Sinha, Neelam Sinha, et al. 694
639 A deep-learning framework to predict cancer treatment re- 695
640 sponse from histopathology images through imputed tran- 696
641 scriptomics. *Nature cancer*, 5(9):1305–1317, 2024. 1 697
- 642 [12] Maximilian Ilse, Jakub Tomczak, and Max Welling. 698
643 Attention-based deep multiple instance learning. pages 699
644 2127–2136, 2018. 2 700
- 645 [13] Guillaume Jaume, Lukas Oldenburg, Anurag Vaidya, 701
646 Richard J Chen, Drew FK Williamson, Thomas Peeters, An- 702
647 drew H Song, and Faisal Mahmood. Transcriptomics-guided 703
slide representation learning in computational pathology. In 704
Proceedings of the IEEE/CVF Conference on Computer Vi- 705
sion and Pattern Recognition, pages 9632–9644, 2024. 2 706
- [14] Hui Liu, Xiaodong Xie, and Bin Wang. Deep learning infers 707
clinically relevant protein levels and drug response in breast 708
cancer from unannotated pathology images. *NPJ Breast* 709
Cancer, 10(1):18, 2024. 1 710
- [15] Titouan Parcollet, Rogier van Dalen, Shucong Zhang, and 711
Sourav Bhattacharya. Summarymixing: A linear-complexity 712
alternative to self-attention for speech recognition and under- 713
standing. *arXiv preprint arXiv:2307.07421*, 2023. 2 714
- [16] Marija Pizurica, Yuanning Zheng, Francisco Carrillo-Perez, 715
Humaira Noor, Wei Yao, Christian Wohlfart, Antoaneta 716
Vladimirova, Kathleen Marchal, and Olivier Gevaert. Digital 717
profiling of gene expression from histology images with 718
linearized attention. *Nature Communications*, 15(1):9886, 719
2024. 2 720
- [17] Benoît Schmauch, Alberto Romagnoni, Elodie Pronier, 721
Charlie Saillard, Pascale Maillé, Julien Calderaro, Aurélie 722
Kamoun, Meriem Sefta, Sylvain Toldo, Mikhail Zaslavskiy, 723
et al. A deep learning model to predict rna-seq expression of 724
tumours from whole slide images. *Nature communications*, 725
11(1):3877, 2020. 1, 2 726
- [18] Charles J Vaske, Stephen C Benz, J Zachary Sanborn, Dent 727
Earl, Christopher Szeto, Jingchun Zhu, David Haussler, and 728
Joshua M Stuart. Inference of patient-specific pathway ac- 729
tivities from multi-dimensional cancer genomics data using 730
paradigm. *Bioinformatics*, 26(12):i237–i245, 2010. 4 731
- [19] Joshua M Wang, Runyu Hong, Elizabeth G Demicco, Jimin 732
Tan, Rossana Lazcano, Andre L Moreira, Yize Li, Anna 733
Calinawan, Narges Razavian, Tobias Schraink, et al. Deep 734
learning integrates histopathology and proteogenomics at a 735
pan-cancer level. *Cell Reports Medicine*, 4(9), 2023. 1 736
- [20] Xiyue Wang, Sen Yang, Jun Zhang, Minghui Wang, 737
Jing Zhang, Wei Yang, Junzhou Huang, and Xiao Han. 738
Transformer-based unsupervised contrastive learning for 739
histopathological image classification. *Medical image anal-* 740
ysis, 81:102559, 2022. 1 741
- [21] Yinxi Wang, Kimmo Kartasalo, Philippe Weitz, Balazs Acs, 742
Masi Valkonen, Christer Larsson, Pekka Ruusuvaori, Johan 743
Hartman, and Mattias Rantalainen. Predicting molecular 744
phenotypes from histopathology images: a transcriptome- 745
wide expression–morphology analysis in breast cancer. *Cancer* 746
research, 81(19):5115–5126, 2021. 1 747

Bridging Morphology and Molecular Signatures: Multi-Task Deep Learning for Multi-Omics Prediction from Histopathology

Supplementary Material

692 7. Top predicted targets

Top Genes	Pearson Correlation	Top Proteins	Pearson Correlation	Top microRNA	Pearson Correlation
AGR3	0.707977	ERALPHA	0.645216	hsa-mir-18a	0.655745
ESR1	0.701862	ASNS	0.604106	hsa-mir-577	0.597472
CCNE1	0.693196	ECADHERIN	0.600070	hsa-mir-17	0.591355
TBC1D9	0.693071	PLK1	0.597685	hsa-mir-942	0.585266
MLPH	0.687592	CYCLINB1	0.586213	hsa-mir-190b	0.579965
THSD4	0.685178	GATA3	0.583716	hsa-mir-210	0.574278
C6orf97	0.684455	CDK1_pT14	0.575290	hsa-mir-505	0.573169
FOXA1	0.682475	FOXM1	0.563957	hsa-mir-590	0.572802
SCUBE2	0.681247	BETACATENIN	0.563820	hsa-mir-130b	0.550578
PSAT1	0.677510	CASPASE7CLEAVEDD198	0.553280	hsa-mir-19a	0.550189
ORC6L	0.676083	IDO	0.545630	hsa-mir-301b	0.548072
CENPA	0.675721	BCL2	0.542713	hsa-mir-301a	0.547313
KCNJ11	0.674130	Aurora-A	0.532306	hsa-mir-934	0.547292
CENPN	0.668963	INPP4B	0.528886	hsa-mir-1307	0.540670
GATA3	0.667477	MSH6	0.527805	hsa-mir-106b	0.538964
FAM54A	0.667091	Mnk1	0.501646	hsa-mir-197	0.535915
DNALI1	0.665293	CYCLINE1	0.500675	hsa-mir-1306	0.535651
CDC25A	0.664781	DJ1	0.488249	hsa-mir-1301	0.533062
DEGS2	0.663996	MCT4	0.488224	hsa-mir-877	0.522733
ZMYND10	0.661833	Puma	0.488050	hsa-mir-16-2	0.519643

Table 1. List of top predicted genes, proteins, microRNA, and pathways.