


Zero-Cost Virtual RNA: Approximating Immunotherapy Signatures via Cross-Modal WSI Retrieval

Sigrid Vila-Bagaria¹ 


SIGRID.VILA@UPC.EDU

Mar Teixidó² 

MTEIXIDO@IRBLLEIDA.CAT

Miquel Piñol² 

MPINOLR.LLEIDA.ICS@GENCAT.CAT

Felip Vilardell² 

FVILARDELL.LLEIDA.ICS@GENCAT.CAT

Robert Montal² 

RMONTAL.LLEIDA.ICS@GENCAT.CAT

Verónica Vilaplana¹ 

VERONICA.VILAPLANA@UPC.EDU

¹ Image Processing Group, Universitat Politècnica de Catalunya, Spain

² Research group of Cancer Biomarkers & Oncological Pathology, IRB Lleida, Spain

Editors: Under Review for MIDL 2026

Abstract

Identifying the “Inflamed” immunophenotype in Gastric Adenocarcinoma predicts immunotherapy response but requires an expensive 10-gene RNA signature. While deep learning on standard H&E slides offers a scalable alternative, conventional binary classifiers oversimplify continuous RNA data and introduce label noise. To resolve this, we propose VITA (VIRtual Transcriptomic Approximation). By aligning H&E and RNA into a joint latent space during training, VITA requires only standard H&E at inference to retrieve morphologically similar historical cases and approximate the continuous RNA signature. Achieving 0.72 classification accuracy and a 0.66 Spearman correlation, VITA provides a cost-effective “virtual transcriptomics” pre-screening tool that preserves the continuous phenotypic spectrum without requiring genomic sequencing.

Keywords: Computational Pathology, Cross-Modal Retrieval, Gastric Cancer.

1. Introduction

Gastric adenocarcinoma (GAC), a leading global cause of cancer mortality (Sung et al., 2021), exhibits marked histological heterogeneity and poor prognosis. Although immune checkpoint inhibitors (ICIs) have transformed GAC treatment (Janjigian et al., 2021, 2023), variable patient response highlights the failure of current biomarkers (Borcoman et al., 2019) to capture complex tumor immunophenotypes. Identifying the “Inflamed” immunophenotype has been proposed to predict ICI response (Rodriguez et al., 2025), yet its 10-gene RNA signature is costly. Deep learning on standard H&E WSIs offers a cost-effective alternative (Lu et al., 2021); however, current WSI classifiers only offer the binary classification, ignoring the biological phenotypic spectrum and discarding molecular nuance.

To overcome this economic and biological trade-off, we propose VIRtual Transcriptomic Approximation (VITA). VITA leverages expensive paired WSI-RNA data during training to learn an aligned metric space, yet requires only standard H&E at inference to eliminate molecular costs. By framing subtyping as a cross-modal retrieval task, VITA searches this latent space for morphological nearest neighbors. This allows pathologists to bypass genomic sequencing and approximate the underlying continuous RNA signature directly from morphology, yielding an interpretable, zero-cost pre-screening tool.

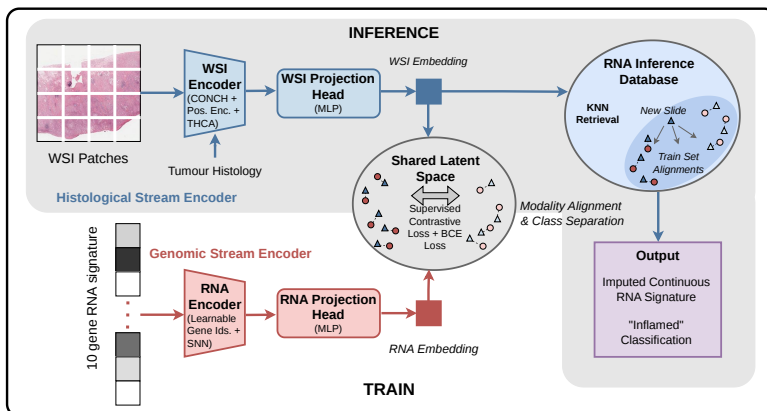


Figure 1: Cross-modal retrieval training and inference framework.

2. Methodology

2.1. Training and Inference

Cross-Modal Architecture and Training: VITA (Fig. 1) employs a dual-stream architecture to learn a continuous joint embedding space. The histological stream extracts visual features via a frozen CONCH backbone (Lu et al., 2024), applying Fourier Positional Encodings and tumour histology-conditioned attention. In parallel, the genomic stream encodes the 10-gene signature using a Self-Normalizing Neural Network (Klambauer et al., 2017). Both modalities are projected into a shared space and optimized jointly using a Supervised Contrastive Loss to transfer genomic knowledge to the visual encoder, alongside a BCE loss to maintain diagnostic separability.

Zero-Cost Inference via Cross-Modal Retrieval At inference, the genomic stream is detached to eliminate molecular costs. Operating unimodally, VITA maps a query H&E slide into the pre-aligned joint space to retrieve its $k = 5$ nearest neighbors, a value determined empirically. Instead of forcing a rigid binary prediction, we aggregate the known RNA profiles of these retrieved cases to impute a continuous RNA signature for the new patient. This search-by-case retrieval framework preserves the spectrum, providing pathologists with an interpretable, zero-cost approximation of the 10-gene RNA signature. Furthermore, while binary classification can be derived directly from this retrieval process, a linear probe is also trained concurrently to perform this specific diagnostic task.

2.2. Experimental Setup

We used a curated cohort of $N = 265$ diagnostic WSIs from the TCGA dataset. Ground-truth labels were derived from the continuous 10-gene RNA signature, yielding 142 “Non-Inflamed” and 123 “Inflamed” cases. Additionally, cases were stratified by pathologists according to the histology Lauren classification (Lauren, 1965) for the histological stream. All models were evaluated using 5-fold cross-validation. To validate our retrieval hypothesis, we compared VITA against CLAM (Lu et al., 2021), the leading unimodal WSI classifier and MCAT (Chen et al., 2021), a standard multimodal fusion network for WSI classification.

Table 1: Performance comparison for immunophenotype subtyping.

Modality	Model	Acc. \uparrow	F1 \uparrow
Unimodal	CLAM	0.73 \pm 0.03	0.72 \pm 0.05
Multimodal	MCAT	0.61 \pm 0.07	0.66 \pm 0.06
Cross-Modal	VITA	0.72 \pm 0.08	0.72 \pm 0.06

Table 2: VITA vs. baseline cross-modal retrieval metrics.

Model	Query \rightarrow Target	R@5 \uparrow	ρ_s \uparrow
CONCH (Avg)	Test \rightarrow Train	–	0.64 \pm 0.04
	Test \rightarrow Test	0.42 \pm 0.05	0.62 \pm 0.08
VITA (Ours)	Test \rightarrow Train	–	0.66 \pm 0.03
	Test \rightarrow Test	0.54 \pm 0.12	0.62 \pm 0.03

3. Results and Discussion

Overcoming the Domain Gap: Direct multimodal fusion (MCAT) struggles in data-scarce regimes (Tab. 1). While unimodal baselines (CLAM) perform well, they rigidly binarize continuous data. VITA addresses this by matching unimodal classification performance while aligning H&E and RNA modalities, setting the foundation for downstream transcriptomic imputation.

Zero-Cost RNA via Retrieval: Operating unimodally at inference, VITA suggests the feasibility of zero-cost pre-screening by querying an H&E slide to retrieve morphological neighbors and averaging their ground-truth signatures into an imputed RNA profile. Emulating a real-world clinical scenario that matches new queries (Test) against a database of known cases (Train), VITA outperforms a baseline of averaged CONCH features (Tab. 2) with superior recall and strong correlation to ground-truth data. Finally, expression barcodes (Fig. 2) validate that VITA effectively reconstructs the phenotypic spectrum without relying on inexact labels.

Limitations & Future Work: While VITA matches strong baselines and correlates well with ground-truth RNA, it remains a proof of concept evaluated on one public cohort, and gains are modest. To establish clinical utility, future work will focus on independent validation and exploring advanced aggregation strategies to enhance imputation fidelity.

4. Conclusions

We present VITA, a framework predicting immunotherapy response in GAC. By aligning WSI-RNA data during training, it enables unimodal H&E inference via cross-modal retrieval. Imputing continuous RNA signatures from morphological neighbors overcomes multimodal data scarcity and label oversimplification. This interpretable approach shows promising potential to bypass genomic sequencing costs and enable zero-cost pre-screening.

Acknowledgments

Supported by AGAUR-FI (2025 FI-STEP 00032) from Generalitat of Catalonia/ESF+, PID2023-148614OB-I00 funded by MICIU/AEI/10.13039/501100011033 and EU FEDER.

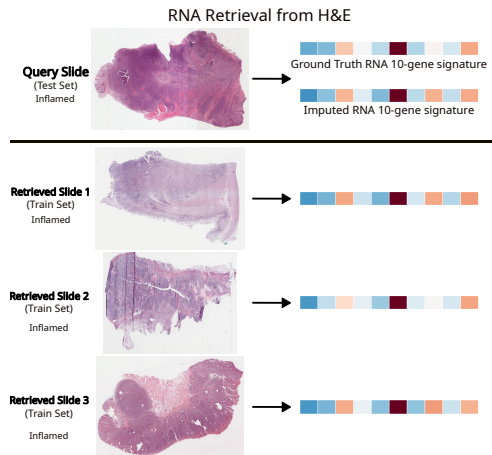


Figure 2: Cross-modal retrieval.

References

- E Borcoman, Y Kanjanapan, S Champiat, S Kato, V Servois, R Kurzrock, S Goel, P Be-
dard, and C Le Tourneau. Novel patterns of response under immunotherapy. *Annals of
Oncology*, 30(3):385–396, 2019.
- Richard J Chen, Ming Y Lu, Wei-Hung Weng, Tiffany Y Chen, Drew FK Williamson,
Trevor Manz, Maha Shady, and Faisal Mahmood. Multimodal co-attention transformer
for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF
International Conference on Computer Vision*, pages 4015–4025, 2021.
- Yelena Y Janjigian, Kohei Shitara, Markus Moehler, Marcelo Garrido, Pamela Salman,
Lin Shen, Lucjan Wyrwicz, Kensei Yamaguchi, Tomasz Skoczykas, Arinilda Campos Bra-
gagnoli, et al. First-line nivolumab plus chemotherapy versus chemotherapy alone for
advanced gastric, gastro-oesophageal junction, and oesophageal adenocarcinoma (check-
mate 649): a randomised, open-label, phase 3 trial. *The Lancet*, 398(10294):27–40, 2021.
- Yelena Y Janjigian, Akihito Kawazoe, Yuxian Bai, Jianming Xu, Sara Lonardi, Jean Phillipe
Metges, Patricio Yanez, Lucjan S Wyrwicz, Lin Shen, Yuriy Ostapenko, et al. Pem-
brolizumab plus trastuzumab and chemotherapy for her2-positive gastric or gastro-
oesophageal junction adenocarcinoma: interim analyses from the phase 3 keynote-811
randomised placebo-controlled trial. *The Lancet*, 402(10418):2197–2208, 2023.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-
normalizing neural networks. *Advances in neural information processing systems*, 30,
2017.
- Pekka Lauren. The two histological main types of gastric carcinoma: diffuse and so-called
intestinal-type carcinoma: an attempt at a histo-clinical classification. *Acta pathologica
microbiologica Scandinavica*, 64(1):31–49, 1965.
- Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri,
and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on
whole-slide images. *Nature biomedical engineering*, 5(6):555–570, 2021.
- Ming Y Lu, Bowen Chen, Drew FK Williamson, Richard J Chen, Ivy Liang, Tong Ding,
Guillaume Jaume, Igor Odintsov, Long Phi Le, Georg Gerber, et al. A visual-language
foundation model for computational pathology. *Nature Medicine*, 30:863–874, 2024.
- Joel Veas Rodriguez, Miquel Piñol, Maria Alba Sorolla, Eva Parisi, Anabel Sorolla, Maria
Santacana, Maria Ruiz, Genís Parra, Mario Bernabeu, Mar Iglesias, et al. Comprehensive
immunophenotyping of gastric adenocarcinoma identifies an inflamed class of tumors
amenable to immunotherapies. *Journal for Immunotherapy of Cancer*, 13(3):e010024,
2025.
- Hyuna Sung, Jacques Ferlay, Rebecca L Siegel, Mathieu Laversanne, Isabelle Soerjo-
mataram, Ahmedin Jemal, and Freddie Bray. Global cancer statistics 2020: Globocan
estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: a
cancer journal for clinicians*, 71(3):209–249, 2021.