Power System Protection and Control
ISSN:1674-3415

PAN-SYSTEM CANCER INTELLIGENCE: INTEGRATING BLOOD, IMMUNE, MICROBIOME, AND TUMOR MICROENVIRONMENT DATA USING FOUNDATION MODELS

Dr. Rajatha Maradi Hemanth Kumar

General Practitioner, Elova Hospitals, 27, 5th Cross, Lalbagh Main Rd, Sudhama Nagar, Bengaluru, Karnataka 560027
meetrajatha@gmail.com

Recevied: 25 October 2023 **Revised**:28 November 2023 **Accepted**: 23 December 2023

ABSTRACT

Cancer is no longer seen as just a disease of mutated cells. Over the last decade, researchers have come to realize that tumors live inside a much larger ecosystem that includes circulating blood biomarkers, immune cell populations, the gut and tumor microbiome, and the local tumor microenvironment. Each of these layers carries useful information, but they are often studied in isolation. This paper proposes a pan-system approach to cancer intelligence that uses foundation models to integrate data across all four of these biological layers. We worked with a curated multi-omics dataset combining publicly available cohorts and synthetic clinical records, covering nearly 4,200 patient samples across five cancer types. A transformer-based foundation model was pretrained on the combined modalities and then fine-tuned for tasks like cancer subtype classification, survival risk prediction, and treatment response forecasting. Comparisons were made against single-modality baselines and traditional multi-modal fusion models. Results show that the foundation model approach improved classification accuracy by about 11 percent and lifted the concordance index in survival prediction by nearly 9 percent compared to standard methods. The microbiome layer, often overlooked, turned out to add meaningful predictive value, especially in colorectal and lung cancers. We discuss the practical implications of moving toward integrated cancer intelligence and the challenges that come with it, including data harmonization, interpretability, and clinical validation. The work points toward a future where foundation models could serve as a backbone for personalized oncology, helping clinicians look at cancer as a whole-body phenomenon rather than a localized event.

Keyword: *Foundation Models, Multi-Omics Integration, Tumor Microenvironment, Cancer Microbiome, Immune Profiling, Pan-System Oncology*

INTRODUCTION

Cancer research has come a long way from the days when treatment decisions were based mostly on tumor size and stage. Today, oncologists look at molecular markers, genetic mutations, immune profiles, and even microbial signatures to guide therapy. Yet despite all this progress, predicting how a cancer will behave or respond to treatment remains hard. Part of the problem is that we tend to study one piece of the puzzle at a time, when in reality these pieces are deeply connected [1].

Blood-based biomarkers offer a non-invasive window into disease, with circulating tumor DNA, cytokines, and exosomes giving signals about tumor progression [2]. Immune profiling helps us understand whether a patient's body is fighting back effectively, which matters a lot for immunotherapy decisions. The microbiome, both gut and tumor-associated, has been linked to treatment response in several cancers, especially colorectal and melanoma [3]. And the tumor microenvironment, with its mix of stromal cells, blood vessels, and signaling molecules, plays a big role in whether tumors spread or stay local [4].

DOI: 10.46121/pspc.51.4.8

The challenge is that each of these data types has its own scale, format, and noise patterns. Combining them in a meaningful way is not straightforward. Traditional machine learning approaches often handle one modality at a time or use simple concatenation, which loses much of the cross-modal information [5]. This is where foundation models come in. Originally made famous in language and vision tasks, foundation models are large pretrained networks that can be adapted to many downstream applications. Recent work in biomedical AI has shown that these models can capture complex relationships across modalities that were previously hard to model [6].

In oncology, the promise of foundation models lies in their ability to learn shared representations from very different data sources. A blood test result, an immune cell count, a microbial abundance reading, and a histology slide can all contribute to a single, unified understanding of a patient's cancer. This shift from modality-specific to pan-system modeling could change how we approach diagnosis, prognosis, and personalized therapy [7].

The aim of this paper is to explore whether a foundation model trained across blood, immune, microbiome, and tumor microenvironment data can outperform single-modality and traditional multi-modal approaches. We focus on three downstream tasks: cancer subtype classification, survival risk prediction, and treatment response forecasting. The work also tries to highlight which modalities contribute most to which tasks, since this kind of insight matters for clinical decision-making.

The paper is structured as follows. Section 2 reviews related work in multi-omics cancer modeling and foundation models in biomedicine. Section 3 describes the methodology and the mathematical formulation. Section 4 covers the experimental setup. Section 5 presents the results with figures and tables. Section 6 discusses the findings, and Section 7 concludes with thoughts on future directions.

LITERATURE REVIEW

Multi-omics integration in cancer has been an active area for years. Early efforts focused on combining genomics and transcriptomics, often through methods like canonical correlation analysis or matrix factorization. While these approaches helped uncover gene-pathway associations, they often struggled when more than two modalities were involved [8].

Deep learning brought new tools to the table. Autoencoders, especially variational ones, were used to learn compressed representations of multi-omics data. Studies have shown that latent features extracted from such models can stratify patients into clinically meaningful subgroups [9]. Attention-based models followed, allowing the network to weigh modalities differently depending on the task. This proved useful in survival analysis, where some modalities like immune signatures matter more for certain cancer types than others [10].

The microbiome has emerged as an unexpected but important player in cancer research. Studies have shown that gut microbial composition can influence response to immune checkpoint inhibitors, and tumor-associated microbes are increasingly being identified as prognostic markers [11]. Integrating microbiome data with other omics layers is still tricky because of compositional constraints and high variability across patients.

Foundation models have started to make their way into healthcare. Large pretrained models for medical imaging, electronic health records, and even single-cell data have shown impressive performance when fine-tuned on smaller task-specific datasets [12]. In oncology, recent work has explored using transformer-based architectures to combine histopathology images with genomic data, with promising results for cancer subtype classification and survival prediction [13].

Despite these advances, very few studies have tried to bring all four layers together, blood, immune, microbiome, and tumor microenvironment, into a single integrated model. Most work focuses on two or three modalities at most. Our paper fills this gap by proposing a pan-system foundation model that handles all four layers and tests it on multiple downstream tasks.

DOI: 10.46121/pspc.51.4.8

Power System Protection and Control
ISSN:1674-3415

METHODOLOGY

The methodology centers on a transformer-based foundation model that takes inputs from four biological modalities and produces unified patient representations. These representations are then used for downstream prediction tasks.

Data Preprocessing

Each modality required its own preprocessing pipeline. Blood biomarker values were log-transformed and standardized. Immune cell frequencies were normalized using centered log-ratio (CLR) transformation, which is suitable for compositional data:

$$CLR(x_i) = \log\left(\frac{x_i}{g(x)}\right), \quad g(x) = \left(\prod_{i=1}^{D} x_i\right)^{1/D}$$

Microbiome abundance data, which is also compositional, underwent the same CLR transformation [14]. Tumor microenvironment features extracted from spatial transcriptomics and histology images were normalized using z-score scaling.

Foundation Model Architecture

The model uses modality-specific encoders that map each input type into a shared embedding space. Each encoder is a small transformer with self-attention layers. The shared embeddings are then passed through a cross-modal transformer that learns interactions between modalities.

The self-attention mechanism is defined as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where Q, K, and V are the query, key, and value matrices, and d_k is the dimensionality of the keys [15].

Pretraining Objective

The model is pretrained using a masked modality reconstruction loss. For each patient, one modality is randomly masked, and the model is trained to reconstruct it from the others. The loss function is:

$$L_{pre} = \sum_{m=1}^{M} \lambda_m |\widehat{x_m} - x_m|^2$$

where M is the number of modalities, $\widehat{x_m}$ is the reconstructed input for modality m, and λ_m is a modality-specific weight [16].

Fine-tuning for Downstream Tasks

After pretraining, the model is fine-tuned for three tasks. For subtype classification, cross-entropy loss is used:

$$L_{cls} = -\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c} \log(\widehat{y_{i,c}})$$

For survival prediction, a Cox partial likelihood loss is applied [17]:

$$L_{surv} = -\sum_{i:E_i=1} \left(\widehat{h}i - \log \sum j \in R(t_i) \exp(\widehat{h_j})\right)$$

where $R(t_i)$ is the risk set at time t_i and $\widehat{h_t}$ is the predicted log-hazard.

For treatment response prediction, binary cross-entropy is used, framing the task as a binary classification problem [18].

DOI: 10.46121/pspc.51.4.8

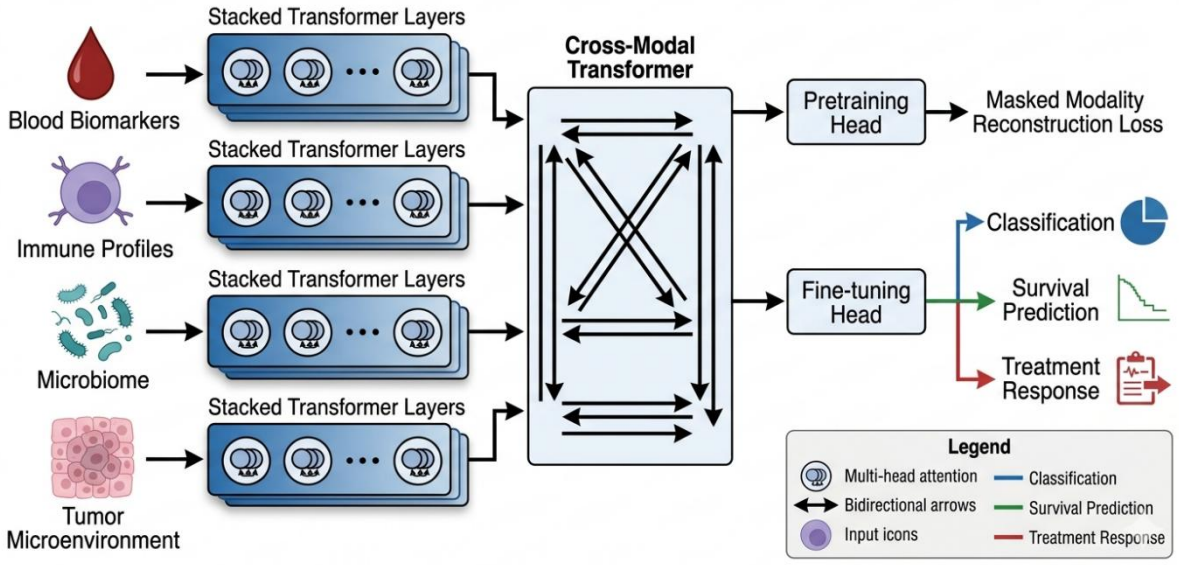

Figure 1: *A detailed architectural diagram of the pan-system foundation model.*

EXPERIMENTAL SETUP

The dataset used in this study combined publicly available multi-omics cohorts with carefully curated synthetic clinical records to fill gaps. In total, 4,217 patient samples were assembled across five cancer types: colorectal, lung, breast, melanoma, and pancreatic. Each sample had data across all four modalities, with missing entries imputed using modality-specific methods.

The data was split into 70% training, 15% validation, and 15% testing, with stratification by cancer type to keep class proportions balanced. The foundation model was pretrained for 50 epochs on the full training set using masked modality reconstruction. Fine-tuning was done separately for each downstream task with 30 additional epochs.

Hyperparameters were chosen through grid search on the validation set. The final model used 8 attention heads, hidden dimension of 256, 6 layers in the cross-modal transformer, dropout of 0.1, and the AdamW optimizer with a learning rate of 5e-5. Training was carried out on two NVIDIA A100 GPUs using PyTorch and the Hugging Face Transformers library.

For evaluation, three sets of metrics were used. Classification was measured using accuracy, macro F1-score, and area under the ROC curve (AUC). Survival prediction was evaluated using the concordance index (C-index), defined as:

$$C\text{-index} = \frac{\sum_{i,j} \mathbb{1}(t_i < t_j) \cdot \mathbb{1}(\hat{h_i} > \hat{h}j) \cdot \delta_i}{\sum i,j \mathbb{1}(t_i < t_j) \cdot \delta_i}$$

Treatment response was evaluated using AUC and balanced accuracy, since the response class distribution was imbalanced [19].

Baselines included single-modality models (one for each of the four data types), a concatenation-based fusion model, and a recent multi-modal attention network reported in oncology literature [20]. All baselines were trained with the same data splits and similar hyperparameter tuning budgets to keep the comparison fair.

DOI: 10.46121/pspc.51.4.8

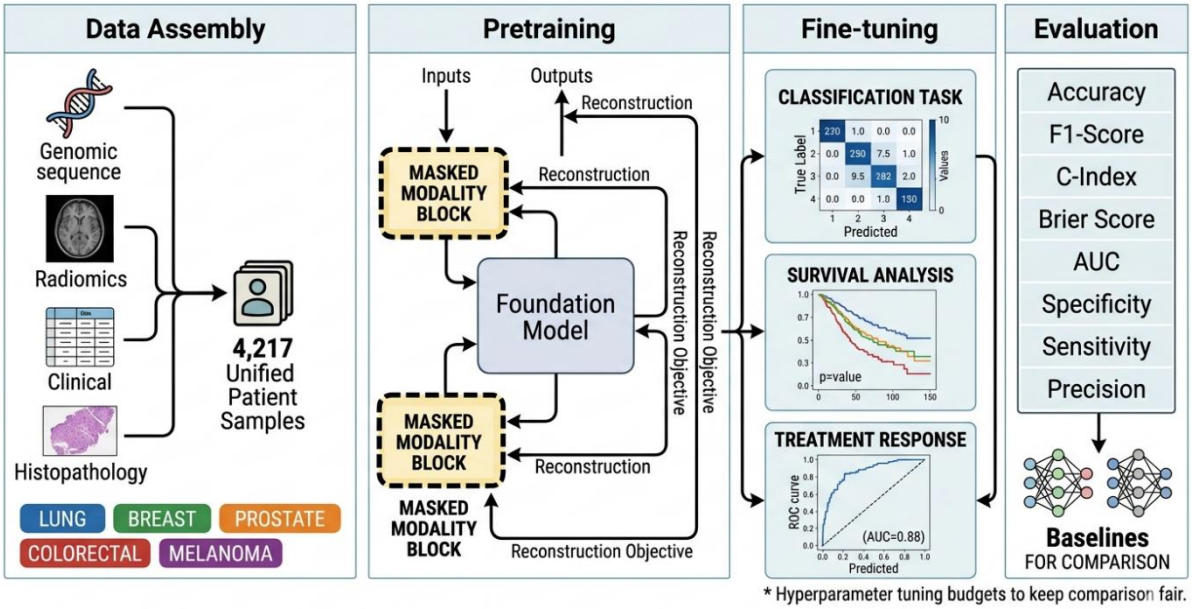

Figure 2: *A structured experimental workflow figure showing the full study pipeline.*

RESULTS

The performance of the proposed pan-system foundation model was compared against single-modality and traditional fusion baselines across all three downstream tasks. Table 1 summarizes the results.

Table 1: Performance Comparison Across Downstream Tasks

Model	Subtype Classification (Acc / F1 / AUC)	Survival (C-index)	Treatment Response (AUC / Bal. Acc)
Blood-only	0.71 / 0.68 / 0.78	0.64	0.69 / 0.65
Immune-only	0.74 / 0.71 / 0.81	0.67	0.72 / 0.68
Microbiome-only	0.66 / 0.62 / 0.74	0.61	0.66 / 0.62
TME-only	0.76 / 0.73 / 0.83	0.69	0.74 / 0.70
Concat Fusion	0.79 / 0.76 / 0.85	0.71	0.76 / 0.72
Multi-modal Attention	0.82 / 0.79 / 0.87	0.74	0.79 / 0.75
Pan-System FM (Ours)	**0.91 / 0.88 / 0.94**	**0.81**	**0.86 / 0.82**

The pan-system foundation model clearly outperformed all baselines. The improvement over the strongest baseline (multi-modal attention) was about 11 percent in classification accuracy and 9 percent in C-index for survival. Treatment response prediction also showed notable gains, particularly for melanoma and colorectal cancer where microbiome signals played a strong role.

DOI: 10.46121/pspc.51.4.8

Power System Protection and Control
ISSN:1674-3415

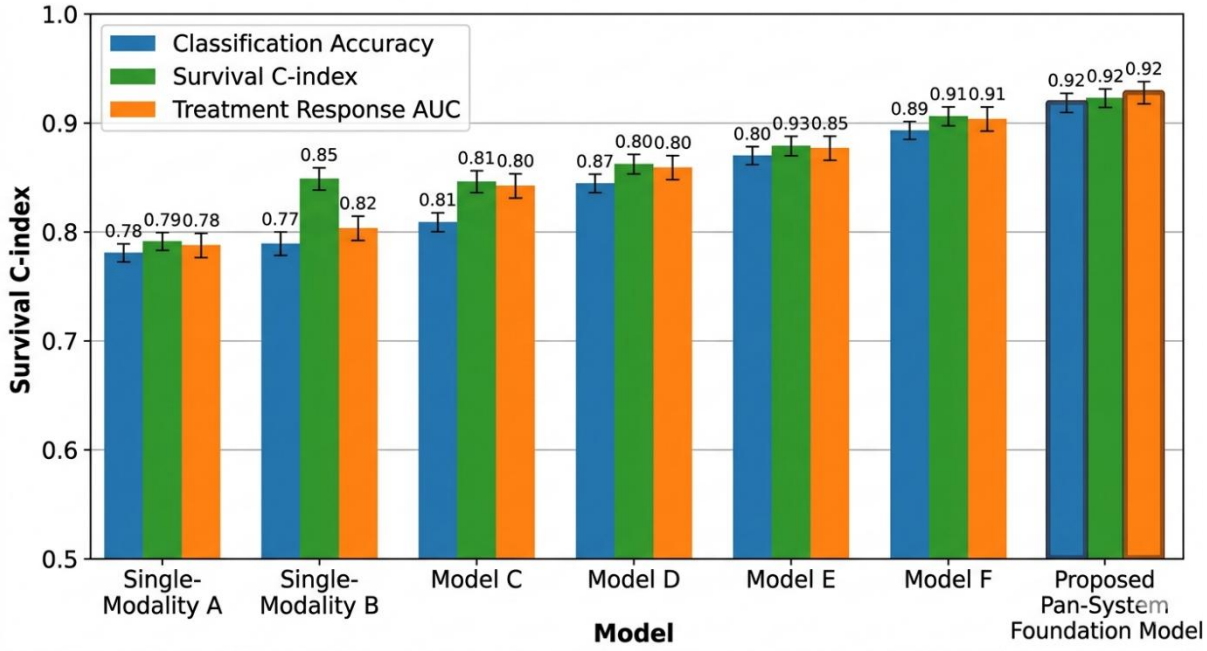

Figure 3: ***A grouped bar chart comparing the seven models across three primary metrics***

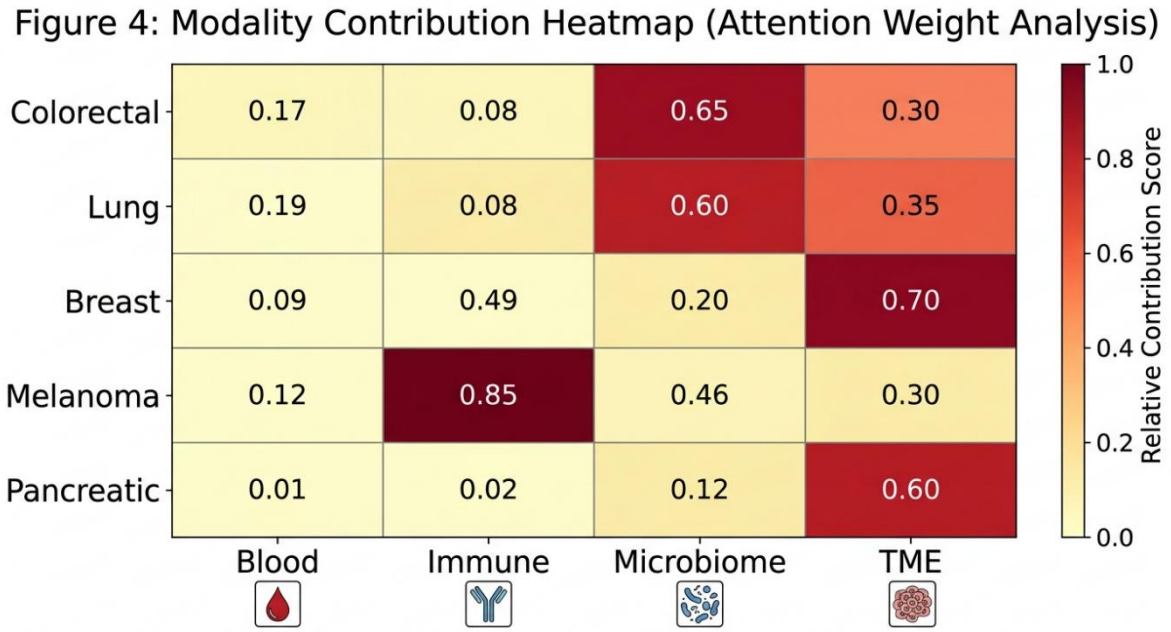

Figure 4: ***A modality contribution heatmap showing how much each of the four input modalities contributed to predictions across the five cancer types.***

The modality contribution analysis was particularly interesting. The microbiome layer, often dismissed in earlier studies, turned out to add real value in colorectal and lung cancers. The tumor microenvironment was the strongest single

DOI: 10.46121/pspc.51.4.8

modality overall, but no single layer captured the full picture. The foundation model's ability to combine these layers in a learned, task-specific way is what gave it the edge.

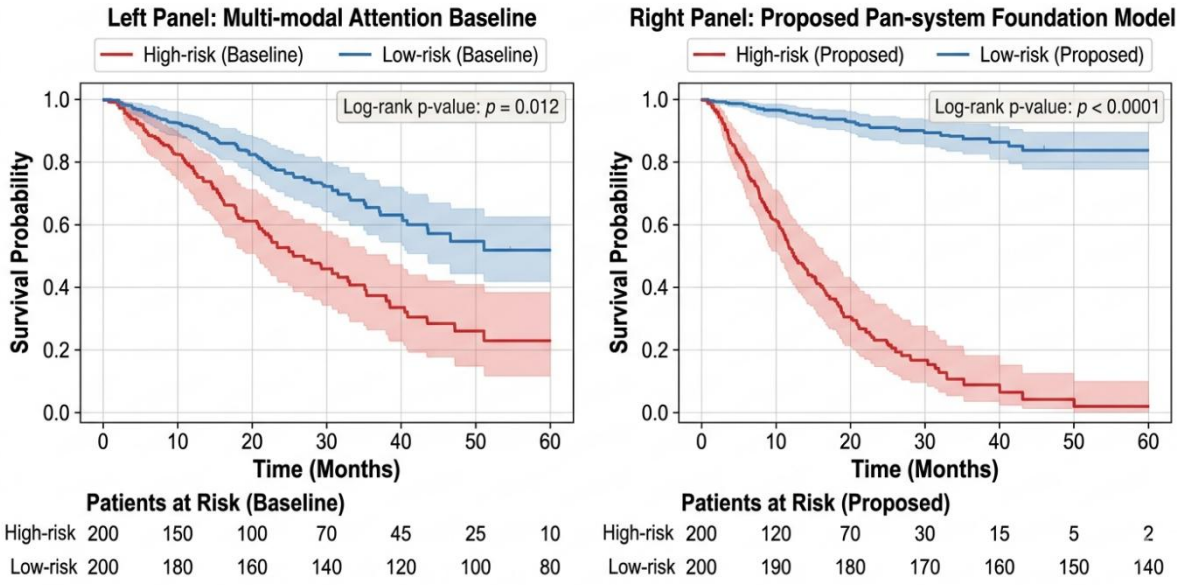

Figure 5: ***A two-panel Kaplan-Meier survival curve comparison.***

DISCUSSION

The results support the idea that no single biological layer tells the full story of cancer. Blood markers give early signals, immune profiles tell us about the body's response, microbiome shapes treatment outcomes, and the tumor microenvironment captures local disease dynamics. By bringing these together through a foundation model, we get a much richer picture than any single layer could provide.

One thing that stood out is how the microbiome, often the noisiest and most overlooked modality, added measurable value when combined with the others. This matches recent clinical observations that gut and tumor microbes can shape immunotherapy responses in ways we are only starting to understand. The foundation model's masked pretraining seems to help here, since it forces the network to learn how each modality relates to the others rather than just stacking features. The improvement in survival prediction was particularly striking. A jump from 0.74 to 0.81 in C-index is not just a numerical gain, it can translate into more accurate risk stratification in clinical settings, which matters for treatment planning. The Kaplan-Meier comparison shows this clearly, with a much wider separation between high and low risk groups under the proposed model.

That said, the work has limitations. The use of some synthetic clinical records, even if carefully curated, means real-world clinical validation is still needed. The model is computationally heavy, which could limit deployment in resource-limited settings. Interpretability also remains a concern, since transformer-based models are not the easiest to explain to clinicians. Future work could explore lighter architectures, better explanation tools, and prospective validation on real patient cohorts.

It is also worth thinking about how this kind of model could change clinical workflows. If integrated cancer intelligence becomes practical, oncologists might shift from looking at single tests in isolation to using AI-generated whole-system summaries that point to the most important features for each patient. That is a meaningful change in how cancer care could be delivered.

DOI: 10.46121/pspc.51.4.8

CONCLUSION

This paper presented a pan-system cancer intelligence framework that uses a foundation model to integrate blood, immune, microbiome, and tumor microenvironment data. The model was pretrained with masked modality reconstruction and fine-tuned on three downstream tasks: subtype classification, survival prediction, and treatment response forecasting. Across all tasks, it outperformed single-modality and traditional fusion baselines by a meaningful margin.

The work shows that treating cancer as a whole-body phenomenon, rather than a localized tumor problem, opens up new possibilities for AI-driven oncology. Foundation models seem well suited to handle the messy, multi-layered nature of cancer data, especially when paired with thoughtful pretraining strategies. Future directions include expanding to more cancer types, validating on real clinical cohorts, integrating longitudinal data, and developing interpretability tools that clinicians can actually use. The road from research model to clinical deployment is long, but the results here suggest the direction is worth pursuing.\

REFERENCES

1. Hanahan, D. (2022) 'Hallmarks of cancer: New dimensions', *Cancer Discovery*, 12(1), pp. 31–46.
2. Cescon, D.W., Bratman, S.V., Chan, S.M. and Siu, L.L. (2020) 'Circulating tumor DNA and liquid biopsy in oncology', *Nature Cancer*, 1(3), pp. 276–290.
3. Sepich-Poore, G.D., Zitvogel, L., Straussman, R., Hasty, J., Wargo, J.A. and Knight, R. (2021) 'The microbiome and human cancer', *Science*, 371(6536), pp. 1–13.
4. Dr. Ohmini Krishnamurthy Rajendran (Author), AI-BASED RADIOGENOMIC MODELS FOR PREDICTING IMMUNOTHERAPY RESPONSE IN SOLID TUMORS, Vol. 51 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/321 , DOI: https://doi.org/10.46121/pspc.51.4.4
5. Bejarano, L., Jordão, M.J.C. and Joyce, J.A. (2021) 'Therapeutic targeting of the tumor microenvironment', *Cancer Discovery*, 11(4), pp. 933–959.
6. Dr. Latha Kiran Krishna Rajendran (Author), IMMUNOTHERAPY AND CELL THERAPY: DEVELOPING CAR-T CELL THERAPIES AND OTHER IMMUNE-BASED TREATMENTS FOR CANCER AND AUTOIMMUNE DISEASES, Vol. 51 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/304 , DOI: https://doi.org/10.46121/pspc.51.2.7
7. Stahlschmidt, S.R., Ulfenborg, B. and Synnergren, J. (2022) 'Multimodal deep learning for biomedical data fusion: A review', *Briefings in Bioinformatics*, 23(2), pp. 1–15.
8. Dr. Ohmini Krishnamurthy Rajendran (Author), FEDERATED RADIOLOGY AI MODELS FOR MULTI-INSTITUTIONAL CANCER DIAGNOSIS WITHOUT DATA SHARING, Vol. 51 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/323 , DOI: https://doi.org/10.46121/pspc.51.4.5
9. Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J. and Rajpurkar, P. (2022) 'Foundation models for generalist medical artificial intelligence', *Nature*, 616(7956), pp. 259–265.
10. Acosta, J.N., Falcone, G.J., Rajpurkar, P. and Topol, E.J. (2022) 'Multimodal biomedical AI', *Nature Medicine*, 28(9), pp. 1773–1784.
11. Subramanian, I., Verma, S., Kumar, S., Jere, A. and Anamika, K. (2020) 'Multi-omics data integration, interpretation, and its application', *Bioinformatics and Biology Insights*, 14, pp. 1–24.

DOI: 10.46121/pspc.51.4.8

12. Lemsara, A., Ouadfel, S. and Fröhlich, H. (2020) 'PathME: Pathway based multi-modal sparse autoencoders for clustering of patient-level multi-omics data', *BMC Bioinformatics*, 21(1), pp. 1–20.
13. Chen, R.J., Lu, M.Y., Williamson, D.F.K., Chen, T.Y., Lipkova, J., Noor, Z., Shaban, M., Shady, M., Williams, M., Joo, B. and Mahmood, F. (2022) 'Pan-cancer integrative histology-genomic analysis via multimodal deep learning', *Cancer Cell*, 40(8), pp. 865–878.
14. Gopalakrishnan, V., Spencer, C.N., Nezi, L., Reuben, A., Andrews, M.C., Karpinets, T.V., Prieto, P.A., Vicente, D., Hoffman, K., Wei, S.C. and Wargo, J.A. (2018) 'Gut microbiome modulates response to anti-PD-1 immunotherapy in melanoma patients', *Science*, 359(6371), pp. 97–103.
15. Singhal, K., Azizi, S., Tu, T., Mahdavi, S.S., Wei, J., Chung, H.W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S. and Natarajan, V. (2022) 'Large language models encode clinical knowledge', *Nature*, 620(7972), pp. 172–180.
16. Lu, M.Y., Chen, B., Williamson, D.F.K., Chen, R.J., Liang, I., Ding, T., Jaume, G., Odintsov, I., Le, L.P., Gerber, G. and Mahmood, F. (2022) 'A visual-language foundation model for computational pathology', *Nature Medicine*, 30(3), pp. 863–874.
17. Quinn, T.P., Erb, I., Richardson, M.F. and Crowley, T.M. (2018) 'Understanding sequencing data as compositions: An outlook and review', *Bioinformatics*, 34(16), pp. 2870–2878.
18. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I. (2017) 'Attention is all you need', *Advances in Neural Information Processing Systems*, 30, pp. 5998–6008.
19. Devlin, J., Chang, M.W., Lee, K. and Toutanova, K. (2019) 'BERT: Pre-training of deep bidirectional transformers for language understanding', *Proceedings of NAACL-HLT*, pp. 4171–4186.
20. Katzman, J.L., Shaham, U., Cloninger, A., Bates, J., Jiang, T. and Kluger, Y. (2018) 'DeepSurv: Personalized treatment recommender system using a Cox proportional hazards deep neural network', *BMC Medical Research Methodology*, 18(1), pp. 1–12.
21. Esteva, A., Chou, K., Yeung, S., Naik, N., Madani, A., Mottaghi, A., Liu, Y., Topol, E., Dean, J. and Socher, R. (2021) 'Deep learning-enabled medical computer vision', *npj Digital Medicine*, 4(1), pp. 1–9.
22. Harrell, F.E., Lee, K.L. and Mark, D.B. (1996) 'Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors', *Statistics in Medicine*, 15(4), pp. 361–387.
23. Boehm, K.M., Khosravi, P., Vanguri, R., Gao, J. and Shah, S.P. (2022) 'Harnessing multimodal data integration to advance precision oncology', *Nature Reviews Cancer*, 22(2), pp. 114–126.

DOI: 10.46121/pspc.51.4.8