Virtual RNA Inference for Survival Prediction in Kidney Renal Clear Cell Carcinoma

Diya Ramakrishnan, Shreyas Pendem
Kyle Coleman, PhD & Joshua Levy, PhD
Emerging Diagnostic and Investigative Technologies, Dartmouth Hitchcock Medical Center

BACKGROUND & MOTIVATION

- Accurate survival prediction for Kidney Renal Clear Cell Carcinoma (KIRC) is necessary for treatment planning and patient outcomes
- ST provides molecular & cellular-level resolution of tumor microenvironment (gene expression, cell types, etc.)
  - Expensive, labor-intensive, rarely available for large clinical cohorts with long-term outcome data.
- Virtual RNA Inference (VRI) can predict spatial gene expression from routine H&E slides
  - H&E slides are low cost and abundant
  - Deepspot (Nonchev, et al.) is a VRI model trained on TCGA (The Cancer Genome Atlas) datasets

OBJECTIVE

OBJECTIVE: evaluate whether predicting gene expression from H&E slides can better forecast kidney cancer patient survival in comparison to just analyzing tissue appearance (morphology) alone.

Secondary Objectives:

- Compare VRI vs. UNI morphological embeddings under identical graph neural network (GNN) architectures
- Assess interpretability through spatial attention maps and UMAP visualization

PRIMARY VS. BASELINE PIPELINE

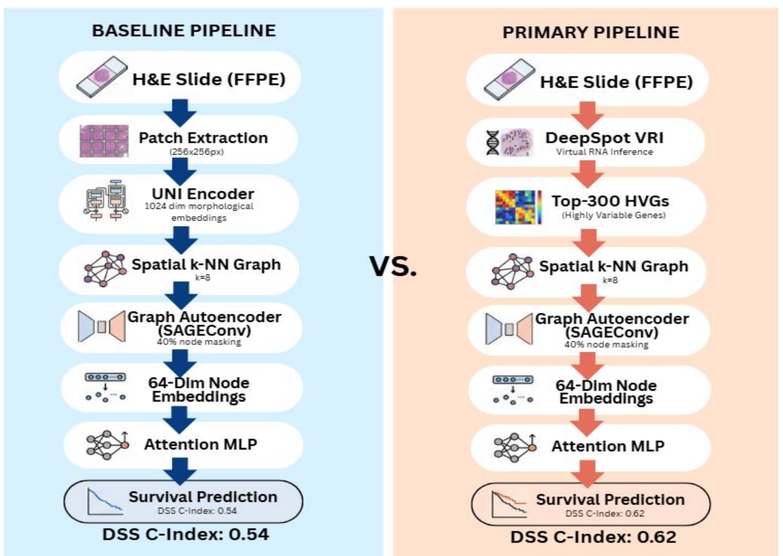

Figure 1. Comparison of 2 pipelines in project (baseline tests survival prediction from tissue appearance alone, while primary tests inferred gene expression for survival prediction

METHODOLOGY

512 TCGA KIRC FFPE patients; 451 training (5-fold CV), 61 held-out test set (Figure 1)

- Baseline: UNI patch embeddings → 1024-dim morphological features per Visium spot
- VRI: DeepSpot-inferred gene expression → top-300 highly variable genes per Visium spot
- Both pipelines: spatial k-NN graph → graph autoencoder → 64-dim embeddings → survival MLP
- Node masking (40%) forces autoencoder to learn tissue context from neighboring spots than memorizing individual features
- Final predictions averaged across all 5 fold models (ensemble)
- Evaluated using Harrell's concordance index (C-index)

RESULTS

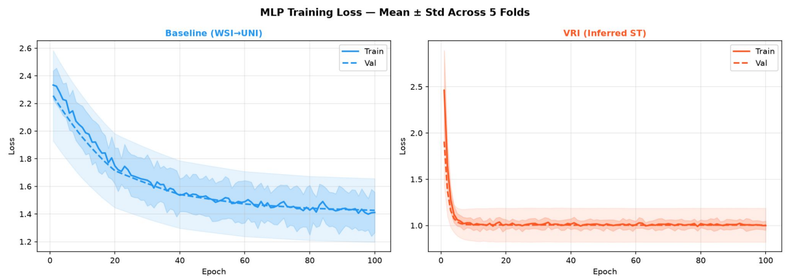

Figure 2. Mean loss curves for both MLPs of both pipelines over five folds

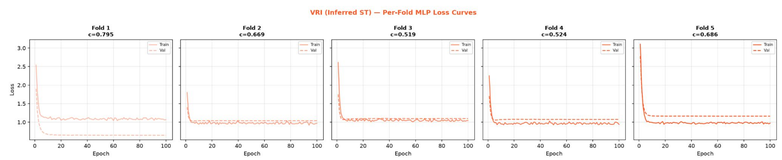

Figure 3. Loss curves for primary pipeline for each of the five folds and c-index

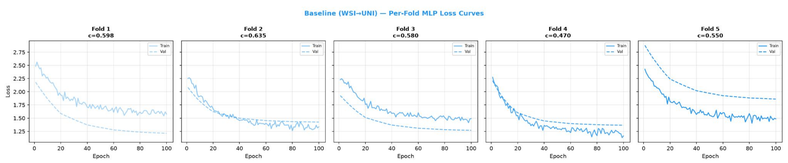

Figure 4. Loss curves for baseline pipeline for each of the five folds and c-index

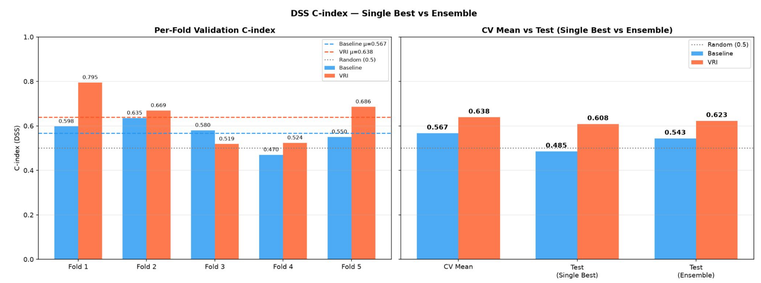

Figure 5. C-indexes for both pipelines over five folds and c-indexes for both pipeline on held-out test set

INTERPRETABILITY

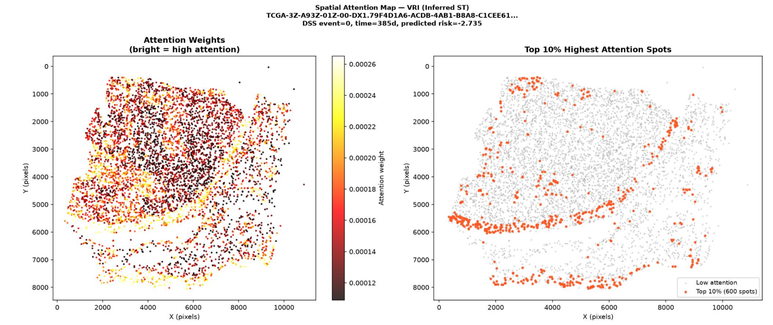

Figure 6. Attention map for VRI; shows Visium spots that model gave most attention

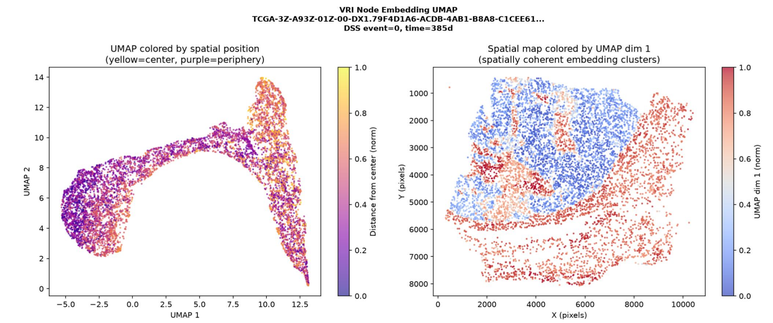

Figure 7. UMAP for VRI node embeddings; shows embedding space by distance from tumor center

DISCUSSION & CONCLUSION

- VRI (C-index 0.622) outperformed the baseline (0.543), suggesting predicted gene expression contains survival information that tissue appearance alone doesn't (Figures 2,3,4,5)
- Model attention concentrated along tumor-stroma boundaries (known prognostic indicator in kidney cancer) (Figure 6)
- VRI embeddings grouped nearby tissue spots together in UMAP without ever being told where spots were located (Figure 7)
- CONCLUSION: models allows for survival prediction from routine H&E slides, applicable to large archives of existing tissue samples lacking gene expression data
- Future work: joint end-to-end training, external KIRC cohort validation, and extension to progression-free interval (PFI)

ACKNOWLEDGMENTS

We'd like to thank the following for their mentorship & support:

- Dr. Joshua Levy & Dartmouth EDIT AI Lab
- Dr. Kyle Coleman from Cedars-Sinai for his mentorship
- Brody McNutt (EDIT AI Team 5 Lead)
- Dartmouth Research Computing, for HPC cluster access
- Dartmouth EDIT AI Lab Summer 2026 Internship Program, for funding and support