



Published in final edited form as:

IEEE Trans Med Imaging. 2022 April ; 41(4): 757–770. doi:10.1109/TMI.2020.3021387.

Pathomic Fusion: An Integrated Framework for Fusing Histopathology and Genomic Features for Cancer Diagnosis and Prognosis

Richard J. Chen,

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA; Broad Institute of Harvard and MIT and the at Dana-Farber Cancer Institute, Boston, MA; Department of Biomedical Informatics, Harvard Medical School, Boston, MA

Ming Y. Lu,

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA; Broad Institute of Harvard and MIT and the at Dana-Farber Cancer Institute, Boston, MA

Jingwen Wang,

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA

Drew F. K. Williamson,

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA

Scott J. Rodig,

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA

Neal I. Lindeman,

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA

Faisal Mahmood

Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA; Broad Institute of Harvard and MIT and the at Dana-Farber Cancer Institute, Boston, MA

Abstract

Cancer diagnosis, prognosis, and therapeutic response predictions are based on morphological information from histology slides and molecular profiles from genomic data. However, most deep learning-based objective outcome prediction and grading paradigms are based on histology or genomics alone and do not make use of the complementary information in an intuitive manner. In this work, we propose *Pathomic Fusion*, an interpretable strategy for end-to-end multimodal fusion of histology image and genomic (mutations, CNV, RNA-Seq) features for survival outcome prediction. Our approach models pairwise feature interactions across modalities by taking the Kronecker product of unimodal feature representations, and controls the expressiveness of each representation via a gating-based attention mechanism. Following supervised learning, we are able to interpret and saliently localize features across each modality, and understand how feature

importance shifts when conditioning on multimodal input. We validate our approach using glioma and clear cell renal cell carcinoma datasets from the Cancer Genome Atlas (TCGA), which contains paired whole-slide image, genotype, and transcriptome data with ground truth survival and histologic grade labels. In a 15-fold cross-validation, our results demonstrate that the proposed multimodal fusion paradigm improves prognostic determinations from ground truth grading and molecular subtyping, as well as unimodal deep networks trained on histology and genomic data alone. The proposed method establishes insight and theory on how to train deep networks on multimodal biomedical data in an intuitive manner, which will be useful for other problems in medicine that seek to combine heterogeneous data streams for understanding diseases and predicting response and resistance to treatment. Code and trained models are made available at: <https://github.com/mahmoodlab/PathomicFusion>.

Keywords

Multimodal Learning; Graph Convolutional Networks; Survival Analysis

I. INTRODUCTION

CANCER diagnosis, prognosis and therapeutic response prediction is usually accomplished using heterogeneous data sources including histology slides, molecular profiles, as well as clinical data such as the patient's age and comorbidities. Histology-based subjective and qualitative analysis of the tumor microenvironment coupled with quantitative examination of genomic assays is the standard-of-care for most cancers in modern clinical settings [1]–[4]. As the field of anatomic pathology migrates from glass slides to digitized whole slide images, there is a critical opportunity for development of algorithmic approaches for joint image-omic assays that make use of phenotypic and genotypic information in an integrative manner.

The tumor microenvironment is a complex milieu of cells that is not limited to only cancer cells, as it also contains immune, stromal, and healthy cells. Though histologic analysis of tissue provides important spatial and morphological information of the tumor microenvironment, the qualitative inspection by human pathologists has been shown to suffer from large inter- and intraobserver variability [5]. Moreover, subjective interpretation of histology slides does not make use of the rich phenotypic information that has been shown to have prognostic relevance [6]. Genomic analysis of tissue biopsies can provide quantitative information on genomic expression and alterations, but cannot precisely isolate tumor-induced genotypic measures and changes from those of non-tumor entities such as normal cells. Current modern sequencing technologies such as single cell sequencing are able to resolve genomic information of individual cells in tumor specimens, with spatial transcriptomics and multiplexed immunofluorescence able to spatially resolve histology tissue and genomics together [7]–[12]. However, these technologies currently lack clinical penetration.

Oncologists often rely on both the qualitative information from histology and quantitative information from genomic data to predict clinical outcomes [13], however, most histology

analysis paradigms do not incorporate genomic information. Moreover, such methods often do not explicitly incorporate information from the spatial organization and community structure of cells, which have known diagnostic and prognostic relevance [6], [14]–[16]. Fusing morphological information from histology and molecular information from genomics provides an exciting possibility to better quantify the tumor microenvironment and harness deep learning for the development of image-omic assays for early diagnosis, prognosis, patient stratification, survival, therapeutic response and resistance prediction.

Contributions:

The contributions of this paper are highlighted as follows:

- **Novel Multimodal Fusion Strategy:** We propose *Pathomic Fusion*, a novel framework for multimodal fusion of histology and genomic features (Fig. 1). Our proposed method models pairwise feature interactions across modalities by taking the Kronecker product of gated feature representations, and controls the expressiveness of each representation using a gating-based attention mechanism.
- **GCNs for Cancer Outcome Prediction:** We present a novel approach for learning cell graph features in histopathology tissue using graph convolutional networks (Fig. 2), and present the first application of GCNs for cancer survival outcome prediction from histology. GCNs act as a complementary method to CNNs for morphological feature extraction, and may be used in lieu of or in combination with CNNs during multimodal fusion for fine-grained patient stratification.
- **Objective Image-Omic Quantitative Study with Multimodal Interpretability:** In a rigorous 15-fold cross-validation-based analysis on two different disease models, we demonstrate that our image-omic fusion paradigm outperforms subjective prognostic determinations that use grading and subtyping, as well as previous state-of-the-art results for patient stratification that use deep learning. To interpret predictions made by our network in survival analysis, we use both class-activation maps and gradient-based attribution techniques to distill prognostic morphological and genomic features.

II. RELATED WORK

Survival Analysis for Cancer Outcome Prediction:

Cancer prognosis via survival outcome prediction is a standard method used for biomarker discovery, stratification of patients into distinct treatment groups, and therapeutic response prediction [17]. With the availability of high-throughput data from next-generation sequencing, statistical survival models have become one of the mainstay approaches for performing retrospective studies in patient cohorts with known cancer outcomes, with common covariates including copy number variation (CNV), mutation status, and RNA sequencing (RNA-Seq) expression [17], [18]. Recent work has incorporated deep learning into survival analysis, in which the covariates for a Cox model are learned using a series of fully connected layers. Yousefi *et al.* [19] proposed using stacked denoising autoencoders to learn a low dimension representation of RNA-Seq data for survival analysis, and in

a follow-up work [19], they used Feedforward Networks to examine the relationship between gene signatures and survival outcomes. Huang *et al.* [20] proposed using weighted gene-expression network analysis as another approach for dimensionality reduction and learning eigen-features from RNA-Seq and micro-RNA data for survival analysis in TCGA. However, these approaches do not incorporate the wealth of multimodal information from heterogeneous data sources including diagnostic slides, which may capture the inherent phenotypic tumor heterogeneity that has known prognostic value.

Multimodal Deep Learning:

Multimodal fusion via deep learning has emerged as an interdisciplinary field that seeks to correlate and combine disparate heterogeneous data modalities to solve difficult prediction tasks in areas such as visual perception, human-computer interaction, and biomedical informatics [21]. Depending on the problem, approaches for multimodal fusion range from fusion of multiview data of the same modality, such as the collection of RGB, depth and infrared measurements for visual scene understanding, to the fusion of heterogeneous data modalities, such as integrating chest X-rays, textual clinical notes, and longitudinal measurements for intensive care monitoring [22]. In the natural language processing community, Kim *et al.* [23] proposed a low-rank feature fusion approach via the Hadamard product for visual question answering, often referred to as bilinear pooling. Zadeh *et al.* [24] studied feature fusion via the Kronecker product for sentiment analysis in audio-visual speech recognition.

Multimodal Fusion of Histology and Genomics:

Though many multimodal fusion strategies have been proposed to address the unique challenges in computer vision and natural language processing, strategies for fusing data in the biomedical domain (e.g. histology images, molecular profiles) are relatively unexplored. In cancer genomics, most works have focused on establishing correspondences between histology tissue and genomics [25]–[28]. For solving supervised learning tasks, previous works have generally relied on the ensembling of extracted feature embeddings from separately trained deep networks (termed late fusion) [20], [29], [30]. Morbadersany *et al.* [29] proposed a strategy for combining histology image and genomic features via vector concatenation. Cheerla *et al.* [31] developed an unsupervised multimodal encoder network for integrating histology image and genomic modalities via concatenation that is resilient to missing data. Shao *et al.* [32] proposed an ordinal multi-modal feature selection approach that identifies important features from both pathological images and multi-modal genomic data, but relies on handcrafted features from cell graph features in histology images. Beyond late fusion, there is limited work in deep learning-based multimodal learning approaches that combine histology and genomic data. Moreover, there is little work made in interpreting histology features in these multimodal deep networks.

Graph-based Histology Analysis:

Though CNNs have achieved remarkable performance in histology image classification and feature representation, graph-based histology analysis has become a promising alternative that rivals many competitive benchmarks. The motivation for interpreting histology images as a graph of cell features (cell graph) is that these computational morphological

(morphometric) features are more easily computed, and explicitly capture cell-to-cell interactions and their spatial organization with respect to the tissue. Prior to deep learning, previous works in learning morphological features from histology images have relied on manually constructed graphs and computing predefined statistics [33]. Doyle *et al.* [34] was the first work to approach Gleason score grading in prostate cancer using Voronoi and Delauney tessellations. Shao *et al.* [32] presented an interesting approach for feature fusion of graph and molecular profile features, with graph features constructed manually and fused via vector concatenation similar to Huang *et al.* [20]. Motivated by the success of representation learning in graphs using deep networks [35]–[38], Anand *et al.* [15], Zhou *et al.*, [39] and Wang *et al.* [16] have used graph convolutional networks for breast, colon and prostate cancer histology classification respectively. Currently, however, there have been no deep learning based-approaches that have used graph convolutional networks for survival outcome prediction.

III. METHODS

Given paired histology and genomic data with known cancer outcomes, our objective is to learn a robust multimodal representation from both modalities that would outperform unimodal representations in supervised learning. Previous works have only relied on CNNs for extracting features from histology images, and late fusion for integrating image features from CNNs with genomic features. In this section, we present our novel approach for integrating histology and genomic data, *Pathomic Fusion*, which fuses histology image, cell graph, and genomic features into a multimodal tensor that explicitly models bimodal and trimodal interactions from each modality. In *Pathomic Fusion*, histology features are extracted as two different views: image-based features using Convolutional Neural Networks (CNNs), and graph-based features using Graph Convolutional Networks (GCNs). Both networks would extract similar morphological features, however, cell graphs from histology images are a more explicit feature representation that directly model cell-to-cell interactions and cell community structure. Following the construction of unimodal features, we propose a gating-based attention mechanism that controls the expressiveness of each feature before constructing the multimodal tensor. The objective of the multimodal tensor is to capture the space of all possible interactions between features across all modalities, with the gating-based attention mechanism used to regularize unimportant features. In subsections A–C, we describe our approach for representation learning in each modality, with subsections D–E describing our multimodal learning paradigm and approach for interpretability. Additional implementation and training details are found in Appendix B.

A. Learning Patient Outcomes from H&E Histology Tissue Images using Convolutional Neural Networks

Anatomic pathology has the ability to reveal the inherent phenotypic intratumoral heterogeneity of cancer, and has been an important tool in cancer prognosis for the past century [40]–[43]. Tumor microenvironment features such as high cellularity and microvascular proliferation have been extensively linked to tumor suppressor deficiency genes and angiogenesis, and recognized to have clinical implications in the recurrence and proliferation of cancer [44]. To capture these features, we train a Convolutional Neural

Network (CNN) on 512×512 image regions-of-interest (ROIs) at $20\times$ magnification ($0.5 \mu\text{m}/\text{pixel}$) as representative regions of cancer pathology. The network architecture of our CNN is VGG19 with batch normalization, which we finetuned using pre-existing weights trained on ImageNet. We extract a $\mathbf{h}_i \in \mathbb{R}^{32 \times 1}$ embedding from the last hidden layer of our Histology CNN, which we use as input into Pathomic Fusion. This network is supervised by the Cox partial likelihood loss for survival outcome prediction, and cross entropy loss for grade classification (Supplement, Appendix A).

B. Learning Morphometric Cell and Graph Features using Graph Convolutional Networks

The spatial heterogeneity of cells in histopathology has potential in informing the invasion and progression of cancer and in bioinformatics tasks of interest such as cancer subtyping, biomarker discovery and survival outcome prediction [40], [45]. Unlike image-based feature representation of histology tissue using CNNs, cell graph representations explicitly capture only pre-selected features of cells, which can be scaled to cover larger regions of histology tissue.

Let $G = (V, E)$ denote a graph with nodes V and edges E . We define $X \in \mathbb{R}^{N \times F}$ as a feature matrix of N nodes in V with F -dimensional features, and $A \in \mathbb{R}^{N \times N}$ as the adjacency matrix that holds the graph topology. To construct graphs that would capture the tumor microenvironment (Fig 2), on the same histology ROI used as input to our CNN, we 1): perform semantic segmentation to detect and spatially localize cells in a histopathology region-of-interest to define our set of nodes V , 2): use K-Nearest Neighbors to find connections between adjacent cells to define our set of edges E , 3): calculate handcrafted and deep features for each cell that would define our feature matrix X , and 4): use graph convolutional networks to learn a robust representation of our entire graph for survival outcome prediction.

Nuclei Segmentation: Accurate nuclei segmentation is important in defining abnormal cell features such as nuclear atypia, abundant tumor cellularity, and other features that would be indicative of cancer progression [46]–[49]. Previous works rely on conventional fully convolutional networks that minimize a pixel-wise loss [50], which can cause the network to segment multiple nuclei as one, leading to inaccurate feature extraction of nuclei shape and community structure. To overcome this issue, we use the same conditional generative adversarial network (cGAN) from our previous work to learn an appropriate loss function for semantic segmentation, which circumvents manually engineered loss functions [51]–[53]. As described in our previous work [51], the conditional GAN framework consists of two networks (a generator G and a discriminator D) that compete against each other in a min-max game to respectively minimize and maximize the objective $\min_G \max_D \mathcal{L}(G, D)$. Specifically, G is a segmentation network that learns to translate histology tissue images n into realistic segmentation masks m , and D is a binary classification network that aims to distinguish real and predicted pairs of tissue ((n, m) vs. $(n, S(n))$). Our generator is supervised with a \mathcal{L}_1 loss and adversarial loss function, in which the adversarial loss penalizes the generator for producing segmentation masks that are unrealistic.

$$\begin{aligned} \mathcal{L}_{\text{GAN}}(S, D_M) = & \mathbb{E}_{m, n \sim p_{\text{data}}(m, n)} [\log D_M(m, n)] \\ & + \mathbb{E}_{n \sim p_{\text{data}}(n)} [\log(1 - D_M(m, S(n)))] \end{aligned}$$

Cell Graph Construction: From our segmented nuclei, we use the K -Nearest Neighbors (KNN) algorithm from the Fast Library for Approximate Nearest Neighbours (FLANN) library to construct the edge set and adjacency matrix of our graph [54] (Fig 2). We hypothesize that adjacent cells will have the most significant cell-cell interactions and limit the adjacency matrix to K nearest neighbours. In our investigations, we used $\mathcal{K} = 5$ to detect community structure and model cellular interactions. Using KNN, our adjacency matrix A is defined as:

$$A_{ij} = \begin{cases} 1 & \text{if } j \in \text{KNN}(i) \text{ and } D(i, j) < d \\ 0 & \text{otherwise} \end{cases}$$

Manual Cell Feature Extraction: For each cell, we computed eight contour features (major axis length, minor axis length, angular orientation, eccentricity, roundness, area, and solidity), as well as four texture features from gray-level co-occurrence matrices (GLCM) (dissimilarity, homogeneity, angular second moment, and energy). Contours were obtained from segmentation results in nuclei segmentation, and GLCMs were calculated from 64×64 image crops centered over each contour centroid. These twelve features were selected for inclusion in our feature matrix X , as they would describe abnormal morphological features about glioma cells such as atypia, nuclear pleomorphism, and hyperchromatism.

Unsupervised Cell Feature Extraction using Contrastive Predictive

Coding: Besides manually computed statistics, we also used an unsupervised technique known as contrastive predictive coding (CPC) [55]–[57] to extract 1024-dimensional features from tissue regions of size 64×64 centered around each cell in a spatial graph. Given a high-dimensional data sequence $\{x_t\}$ (256×256 image crop from the histology ROI), CPC is designed to capture high-level representations shared among different portions (64×64 image patches) of the complete signal. The encoder network g_{enc} transforms each data observation x_t into a low-dimensional representation z_t and learns via a contrastive loss whose optimization leads to maximizing the mutual information between the available context c_t , computed from a known portion of the encoded sequence $\{z_j\}$, $j < t$ and future observations z_{t+k} , $k > 0$. By minimizing the CPC objective, we are able to learn rich feature representations shared among various tissue regions that are specific to the cells in the underlying tissue site. Examples include the morphology and distinct arrangement of different cell types, inter-cellular interactions, and the microvascular patterns surrounding each cell. To create CPC features for each cell, we encode 64×64 image patches centered over the centroid of each cell. These features are concatenated with our handcrafted features during cell graph construction.

Graph Convolutional Network: Similar to CNNs, GCNs learn abstracts feature representations for each feature in a node via message passing, in which nodes iteratively aggregate feature vectors from their neighborhood to compute a new feature vector at the

next hidden layer in the network [38]. The representation of an entire graph can be obtained through pooling over all the nodes, which can then be used as input for tasks such as classification or survival outcome prediction. Such convolution and pooling operations can be defined as follows:

$$\begin{aligned} a_v^{(k)} &= \text{AGGREGATE}^{(k)}(\{h_u^{(k-1)}: u \in \mathcal{N}(v)\}) \\ h_v^{(k)} &= \text{COMBINE}^{(k)}(h_v^{(k-1)}, a_v^{(k)}) \end{aligned}$$

where $h_v^{(k)}$ is the feature vector of node v at the $k-1$ -th iteration of the neighborhood aggregation, $a_v^{(k)}$ is the feature vector of node v at the next iteration, and AGGREGATE and COMBINE are functions for combining feature vectors between hidden layers. As defined in Hamilton *et al.*, we adopt the AGGREGATE and COMBINE definitions from GraphSAGE [35], which for a given node, represents the next node hidden layer as the concatenation of the current hidden layer with the neighborhood features:

$$\begin{aligned} a_v^{(k)} &= \text{MAX}(\{\text{ReLU}(W \cdot h_u^{(k-1)}), \forall u \in \mathcal{N}(v)\}) \\ h_v^{(k)} &= W \cdot [h_v^{(k-1)}, a_v^{(k)}] \end{aligned}$$

Unlike other graph-structured data, cell graphs exhibit a hierarchical topology, in which the degree of eccentricity and clustered components of nodes in a graph define multiple views of how cells are organized in the tumor micro-environment: from fine-grained views such as local cell-to-cell interactions, to coarser-grained views such as structural regions of cell invasion and metastasis. In order to encode the hierarchical structure of cell graphs, we adopt the self-attention pooling strategy SAGPOOL presented in Lee *et al.* [36], which is a hierarchical pooling method that performs local pooling operations of node embeddings in a graph. In attention pooling, the contribution of each node embedding in the pooling receptive field to the next network layer is adaptively learned using an attention mechanism. The attention score $Z \in \mathbb{R}^{N \times 1}$ for nodes in G can be calculated as such:

$$Z = \sigma(\text{SAGEConv}(X, A + A^2))$$

where X are the node features, A is the adjacency matrix, and SAGEConv is the convolution operator from GraphSAGE. To also aggregate information from multiple scales in the nuclei graph topology, we also adopt the hierarchical pooling strategy in Lee *et al.* [36]. Since we are constructing cell graphs on the entire image, no patch averaging of predicted hazards needs to be performed. At the last hidden layer of our Graph Convolutional SNN, we pool the node features into a $\mathbf{h}_g \in \mathbb{R}^{32 \times 1}$ feature vector, which we use as an input to Pathomic Fusion.

C. Predicting Patient Outcomes from Molecular Profiles using Self-Normalizing Networks

Advances in next-generation sequencing data have allowed for the profiling of transcript abundance (RNA-Seq), copy number variation (CNV), mutation status, and other molecular characterizations at the gene level, and have been frequently used to study survival outcomes in cancer. For example, isocitrate dehydrogenase 1 (IDH1) is a gene that is important for

cellular metabolism, epigenetic regulation and DNA repair, with its mutation associated with prolonged patient survival in cancers such as glioma. Other genes include EGFR, VEGF and MGMT, which are implicated in angiogenesis, which is the process of blood vessel formation that also allows cancer to proliferate to other areas of tissue.

For learning scenarios that have hundreds to thousands of features with relatively few training samples, Feedforward networks are prone to overfitting. Compared to other kinds of neural network architectures such as CNNs, weights in Feedforward networks are shared and thus more sensitive training instabilities from perturbation and regularization techniques such as stochastic gradient descent and Dropout. To mitigate overfitting on high-dimensional low sample size genomics data and employ more robust regularization techniques when training Feedforward networks, we adopt the normalization layers from Self-Normalizing Networks in Klambauer *et al.* [58]. In Self-Normalizing Networks (SNN), rectified linear unit (ReLU) activations are replaced with scaled exponential linear units (SeLU) to drive outputs after every layer towards zero mean and unit variance. Combined with a modified regularization technique (Alpha Dropout) that maintains this self-normalizing property, we are able to train well-regularized Feedforward networks that would be otherwise prone to instabilities as a result of vanishing or explosive gradients. Our network architecture consists of four fully-connected layers followed by Exponential Linear Unit (ELU) activation and Alpha Dropout to ensure the self-normalization property. The last fully-connected layer is used to learn a representation $\mathbf{h}_n \in \mathbb{R}^{32 \times 1}$, which is used as input into our Pathomic Fusion (Fig. 1).

D. Multimodal Tensor Fusion via Kronecker Product and Gating-Based Attention

For multimodal data in cancer pathology, there exists a data heterogeneity gap in combining histology and genomic input - histology images are spatial distributed as (R, G, B) pixels in a two-dimensional grid, whereas cell graphs are defined as a set of nodes V with different sized neighborhoods and edges E , and genomic data is often represented as a one-dimensional vector of covariates [30]. Our motivation for multimodal learning is that the inter-modality interactions between histology and genomic features would be able to improve patient stratification into subtypes and treatment groups. For example, in the refinement of histogenesis of glioma, though morphological characteristics alone do not correlate well with patient outcomes, their semantic importance in drawing decision boundaries is changed when conditioned on genomic biomarkers such as IDH1 mutation status and chromosomal 1p19q arm codeletion [59].

In this work, we aim to explicitly capture these important interactions using the Kronecker Product, which model feature interactions across unimodal feature representations, that would otherwise not be explicitly captured in feedforward layers. Following the construction of the three unimodal feature representations in the previous subsections, we build a multimodal representation using the Kronecker product of the histology image, cell graph, and genomic features $(\mathbf{h}_i, \mathbf{h}_g, \mathbf{h}_n)$. The joint multimodal tensor computed by the matrix outer product of these feature vectors would capture important unimodal, bimodal and trimodal interactions of all features of these three modalities, shown in Fig. 1 and in the equation below:

$$\mathbf{h}_{\text{fusion}} = \begin{bmatrix} \mathbf{h}_i \\ 1 \end{bmatrix} \otimes \begin{bmatrix} \mathbf{h}_g \\ 1 \end{bmatrix} \otimes \begin{bmatrix} \mathbf{h}_n \\ 1 \end{bmatrix}$$

where \otimes is the outer product, and $\mathbf{h}_{\text{fusion}}$ is a differential multimodal tensor that forms in a 3D Cartesian space. In this computation, every neuron in the last hidden layer in the CNN is multiplied by every other neuron in the last hidden layer of the SNN, and subsequently multiplied with every other neuron in the last hidden layer of the GCN. To preserve unimodal and bimodal feature interactions when computing the trimodal interactions, we append 1 to each unimodal feature representation. For feature vectors of size $[33 \times 1]$, $[33 \times 1]$ and $[33 \times 1]$, the calculated multimodal tensor would have dimension $[33 \times 33 \times 33]$, where the unimodal features (\mathbf{h}_i , \mathbf{h}_g , \mathbf{h}_n) and bimodal feature interactions ($\mathbf{h}_i \otimes \mathbf{h}_g$, $\mathbf{h}_g \otimes \mathbf{h}_n$, $\mathbf{h}_i \otimes \mathbf{h}_n$) are defined along the outer dimension of the 3D tensor, and the trimodal interactions (captured as $\mathbf{h}_i \otimes \mathbf{h}_g \otimes \mathbf{h}_n$) in the inner dimension of the 3D tensor (Fig. 1). Following the computation of this joint representation, we learn a final network using fully-connected layers using the multimodal tensor as input, supervised with the previously defined Cox objective for survival outcome prediction and cross-entropy loss for grade classification. Ultimately, the value of Pathomic Fusion is fusing heterogeneous modalities that have disparate structural dependencies. Our multimodal network is initialized with pretrained weights from the unimodal networks, followed by end-to-end fine-tuning of the Histology GCN and Genomic SNN.

To decrease the impact of noisy unimodal features during multimodal training, before the Kronecker Product, we employed a gating-based attention mechanism that controls the expressiveness of features of each modality [60]. In fusing histology image, cell graph, and genomic features, some of the captured features may have high collinearity, in which employing a gating mechanism can reduce the size of the feature space before computing the Kronecker Product. For a modality m with a unimodal feature representation \mathbf{h}_m , we learn a linear transformation $W_{ign \rightarrow m}$ of modalities \mathbf{h}_i , \mathbf{h}_g , \mathbf{h}_n that would score the relative importance of each feature in m , denoted as \mathbf{z}_m in the equation below.

$$\begin{aligned} \mathbf{h}_{m, \text{gated}} &= \mathbf{z}_m * \mathbf{h}_m, \forall m \in \{i, g, n\} \\ \text{where, } \mathbf{h}_m &= \text{ReLU}(W_m \cdot \mathbf{h}_m) \\ \mathbf{z}_m &= \sigma(W_{ign \rightarrow m} \cdot [\mathbf{h}_i, \mathbf{h}_g, \mathbf{h}_n]) \end{aligned}$$

\mathbf{z}_m can be interpreted as an attention weight vector, in which modalities i , g , n attend over each feature in modality m . W_m and $W_{ign \rightarrow m}$ are weight matrix parameters we learn for feature gating. After taking the softmax probability, we take the element-wise product of features \mathbf{h}_m and scores \mathbf{z}_m to calculate the gated representation.

E. Multimodal Interpretability

To interpret our network, we modified both Grad-CAM and Integrated Gradients for visualizing image saliency feature importance across multiple types of input. Grad-CAM is a gradient-based localization technique used to produce visual explanations in image classification, in which neurons whose gradients have positive influence on a class of

interest are used to produce a coarse heatmap [61]. Since the last layer of our network is a single neuron for outputting hazard, we modified the target to perform back-propagation on the single neuron. As a result, the visual explanations from our network correspond with image regions used in predicting hazard (values ranging from $[-3,3]$). For Histology GCN and Genomic SNN, we used Integrated Gradients (IG), a gradient-based feature attribution method that attributes the prediction of deep networks to their inputs [62]. Similar to previous attribution-based methods such as Layer-wise Relevance Propagation [63], IG calculates the gradients of the input tensor x across different scales against a baseline x_j (zero-scaled), and then uses the Gauss-Legendre quadrature to approximate the integral of gradients.

$$\text{IG}_i(x) := (x_i - x'_i) \times \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha \times (x - x'))}{\partial x_i} d\alpha$$

To adapt IG to graph-based structures, we treat the nodes in our graph input as the batch dimension, and scale each node in the graph by the number of integral approximation steps. With multimodal inputs, we can approximate the integral of gradients for each data modality.

IV. EXPERIMENTAL SETUP

A. Data Description

To validate our proposed multimodal paradigm for integrating histology and genomic features, we collected glioma and clear cell renal cell carcinoma data from the TCGA, a cancer data consortium that contains paired high-throughput genome analysis and diagnostic whole slide images with ground-truth survival outcome and histologic grade labels. For astrocytomas and glioblastomas in the merged TCGA-GBM and TCGA-LGG (TCGA-GBMLGG) project, we used 1024×1024 region-of-interests (ROIs) from diagnostic slides curated by [29], and used sparse stain normalization [64] to match all images to a standard H&E histology image. Multiple region-of-interests (ROIs) from diagnostic slides were obtained for some patients, creating a total of 1505 images for 769 patients. 320 genomic features from CNV (79), mutation status (1), and bulk RNA-Seq expression from the top 240 differentially expressed genes (240) were curated from the TCGA and the cBioPortal [65] for each patient. For clear cell renal cell carcinoma in the TCGA-KIRC project we used manually extracted 512×512 ROIs from diagnostic whole slide images. For 417 patients in CCRCC, we collected 3 512×512 40x ROIs per patient, yielding 1251 images total that were similarly normalized with stain normalization. We paired these images with 357 genomic features from CNV of genes with alteration frequency greater than 7% (117) and RNA-Seq from the top 240 differentially expressed genes (240). It should be noted that for TCGA-GBMLGG had approximately 40% of the patients had missing RNA-Seq expression. Details regarding genomic features and data alignment of histology and genomics data are found in the implementation details (Appendix B). Our experimental setup is also described in the reproducibility section of our GitHub repository.

B. Quantitative Study

TCGA-GBMLGG: Gliomas are a form of brain and spinal cord tumors defined by both hallmark histopathological and genomic heterogeneity in the tumor microenvironment, as well as response-to-treatment heterogeneity in patient outcomes. The current World Health Organization (WHO) Paradigm for glioma classification stratifies diffuse gliomas based on morphological and molecular characteristics: glial cell type (astrocytoma, oligodendroglioma), IDH1 gene mutation status and 1p19q chromosome codeletion status [59]. WHO Grading is made by the manual interpretation of histology using pathological determinants for malignancy (WHO Grades II, III, and IV). These characteristics form three categories of gliomas which have been extensively correlated with survival: 1) IDH-wildtype astrocytomas (IDHwt ATC), 2) IDH-mutant astrocytomas (IDHmut ATC), and 3) IDH-mutant and 1p/19q-codeleted oligodendrogliomas (ODG). IDHwt ATCs (predominantly WHO grades III and IV) have been shown to have the worst patient survival outcomes, while IDHmut ATCs (mixture of WHO Grades II, III, and IV) and ODGs (predominantly WHO Grades II and III) have more favorable outcomes (listed in increasing order) [59]. As a baseline against standard statistical approaches / WHO paradigm for survival outcome prediction, we trained Cox Proportion Hazard Models using age, gender, molecular subtypes and grade as covariates.

In our experimentation, we conducted an ablation study comparing model configurations and fusion strategies in a 15-fold cross validation on two supervised learning tasks for glioma: 1) survival outcome prediction, and 2) cancer grade classification. For each task, we trained six different model configurations from the combination of available modalities in the dataset. First, we trained three different unimodal networks: 1) a CNN for in histology image input (Histology CNN), 2) a GCN for cell graph input (Histology GCN), and 3) a SNN for genomic features input (Genomic SNN). For cancer grade classification, we did not use mRNA-Seq expression due to missing data, lack of paired training examples, and because grade is solely determined from histopathologic appearance. After training the unimodal networks, we trained three different configurations of Pathomic Fusion: 1) GCN \otimes SNN, 2) CNN \otimes SNN, 3) GCN \otimes CNN \otimes SNN. To test for ensembling, we train multimodal networks that fused histology data with histology data, and genomic features with genomic features. We compare our fusion approach to internal benchmarks and the previous state-of-the-art [29] approach for survival outcome prediction in glioma, which concatenates histology ROIs with IDH1 and 1p19q genomic features. To compare with their results, we used their identical train-test split, which was created using a 15-fold Monte Carlo cross-validation [29].

TCGA-KIRC: Clear cell renal cell carcinoma (CCRCC) is the most common type of renal cell carcinoma, originating from cells in the proximal convoluted tubules. Histopathologically, CCRCC is characterized by diverse cystic grown patterns of cells with clear or eosinophilic cytoplasm, and a network of thin-walled “chicken wire” vasculature [66], [67]. Genetically, it is characterized by a chromosome 3p arm loss and mutation status of the von Hippel-Lindau (VHL) gene, which leads to lead to stabilization of hypoxia inducible factors that lead to malignancy [68]. Though CCRCC is well-characterized, methods for staging CCRCC suffer from large intra-observer variability

in visual histopathological examination. The Fuhrman Grading System for CCRCC is a nuclear grade that ranges from G1 (round or uniform nuclei with absent nucleoli) to G4 (irregular and multilobular nuclei with prominent nucleoli). At the time of the study, the TCGA-KIRC project used the Fuhrman Grading System to grade CCRCC in severity from G1 to G4, however, the grading system has received scrutiny in having poor overall agreement amongst pathologists on external cohorts [66]. As a baseline against standard statistical approaches, we trained Cox Proportion Hazard Models using age, gender, and grade as covariates.

Similar to the ablation study conducted with glioma, we compared model configurations and fusion strategies in a 15-fold cross validation on CCRCC, and tested for ensembling effects. In demonstrating the effectiveness of Pathomic Fusion in stratifying CCRCC, we use the Fuhrman Grade as a comparative baseline in survival analysis, however, we do not perform ablation experiments on grade classification. Since CCRCC does not have multiple molecular subtypes, subtyping was also not performed, however, we perform analyses on CCRCC patient cohorts with different survival durations (shorter surviving and longer surviving patients).

Evaluation: We evaluate our method with standard quantitative and statistical metrics for survival outcome prediction and grade classification. For survival analysis, we evaluate all models using the Concordance Index (c-Index), which is defined as the fraction of all pairs of samples whose predicted survival times are correctly ordered among all uncensored samples (Table I, II). On glioma and CCRCC respectively, we separate the predicted hazards into 33-66-100 and 25-50-75-100 percentiles as digital grades, which we compared with molecular subtyping and grading. For significance testing of patient stratification, we use the Log Rank Test to measure if the difference of two survival curves is statistically significant [69]. Kaplan-Meier estimates and predicted hazard distribution were used to visualize how models were stratifying patients. For grade classification, we evaluate our networks using Area Under the Curve (AUC), Average Precision (AP), F1-Score (micro-averaged across all classes), F1-Score (WHO Grade IV class only), and show ROC curves (Appendix C, Fig. 7). In total, we trained 480 models total in our ablation experiments using 15-fold cross validation. Implementation and training details for all networks are described in detail in Appendix A and B.

V. RESULTS AND DISCUSSION

A. Pathomic Fusion Outperforms Unimodal Networks and the WHO Paradigm

In combining histology image, cell graph, and genomic features via Pathomic Fusion, our approach outperforms Cox models, unimodal networks, and previous deep learning-based feature fusion approaches on image-omic-based survival outcome prediction (Table I, II). On glioma, Pathomic Fusion outperforms the WHO paradigm and the previous state-of-the-art (concatenation-based fusion [29]) with **6.31%** and **5.76%** improvements respectively, reaching a c-Index of **0.826**. In addition, we demonstrate that multimodal networks were able to consistently improve upon their unimodal baselines, with trimodal Pathomic Fusion (CNN \otimes GCN \otimes SNN) fusion of image, graph, and genomic features having the largest

c-Index. Though bimodal Pathomic Fusion (CNN \otimes SNN) achieved similar performance metrics, the difference between low-to-intermediate digital grades ([0,33] vs. [33,66] percentile of predicted hazards) was not found to be statistically significant (Appendix C, Table III, Fig. 9). In incorporating features from GCNs, the p-value for testing difference in [0,33] vs. [33,66] percentiles decreased from 0.103 to 2.68e-03. On CCRCC, we report similar observations, with trimodal Pathomic Fusion achieving a c-Index of **0.720** and statistical significance in stratifying patients into low and high risk (Appendix C, Fig. 10). Using the c-Index metric, GCNs do not add significant improvement over CNNs alone. However, for heterogeneous cancers such as glioma, the integration of GCNs in Pathomic Fusion may provide clinical benefit in distinguishing survival curves of less aggressive tumors.

We also demonstrate that these improvements are not due to network ensembling, as inputting same modality twice into Pathomic Fusion resulted in overfitting (Appendix C, Table III). On glioma grade classification, we see similar improvement with Pathomic Fusion with increases of 2.75% AUC, 4.23% average precision, 4.27% F1-score (micro), and 5.11% (Grade IV) over Histology CNN, which is consistent with performance increases found in multimodal learning literature for conventional vision tasks (Appendix C, Fig. 8, Table V) [24].

B. Pathomic Fusion Improves Patient Stratification

To further investigate the ability of Pathomic Fusion for improving objective image-omic-based patient stratification, we plot Kaplan-Meier (KM) curves of our trained networks against the WHO paradigm (which uses molecular subtyping) on glioma (Fig. 3), and against the Fuhrman Grading System on CCRCC (Fig. 4). Overall, we observe that Pathomic Fusion allows for fine-grained stratification of survival curves beyond low vs. high survival, and that these digital grades may be useful in clinical settings in defining treatment cohorts.

On glioma, similar to [29], we observe that digital grading (33-66 percentile) from Pathomic Fusion is similar to that of the three defined glioma subtypes (IDHwt ATC, IDHmut ATC, ODG) that correlate with survival. In comparing Pathomic Fusion to Histology CNN, Pathomic Fusion was able to discriminate intermediate and high risk patients better than Histology CNN. Though Pathomic Fusion was slightly worse in defining low and intermediate risk patients, differences between these survival curves were observed to be statistically significant (Appendix C, Table III). Similar confusion in discriminating low-to-intermediate risk patients is also shown in the KM estimates of molecular subtypes, which corroborates with known literature that WHO Grades II and III are more difficult to distinguish than Grades III and IV [59] (Fig. 3). In analyzing the distribution of predicted hazard scores for patients in low vs. high surviving cohorts, we also observe that Pathomic Fusion is able to correctly assign risk to these patients in three high-density peaks / clusters, whereas Histology CNN alone labels a majority of intermediate-to-high risk gliomas with low hazard values. In inspecting the clusters elucidated by Pathomic Fusion ($[-1.0, -0.5]$, $[1.0, 1.25]$ and $[1.25, 1.5]$), we see that the gliomas these clusters strongly corroborate with the WHO Paradigm for stratifying gliomas into IDHwt ATC, IDHmut ATC, and ODG.

On CCRCC, we observe that Pathomic Fusion is able to not only differentiate between lower and higher surviving patients, but also assign digital grades that follow patient stratification by the Fuhrman Grading System (Fig. 4). Unlike Histology CNN, Pathomic Fusion is able to disentangle the survival curves of G1-G3 CCRCCs, which have overall low-to-intermediate survival. In analyzing the distribution of hazard predictions by Histology CNN, we see that risk is almost uniformly predicted across shorter and longer survival patients, which suggests that histology alone is a poor prognostic indicator for survival in CCRCC.

C. Multimodal Interpretability of Pathomic Fusion

In addition to improved patient stratification, we demonstrate that our image-omic paradigm is highly interpretable, in which we can attribute how pixel regions in histology images, cells in cell graphs, and features in genomic inputs are used in survival outcome prediction.

In examining IG attributions for genomic input, we were able to corroborate important markers such as IDH wildtype status in glioma and CYP3A7 under-expression in CCRCC correlating with increased risk. In glioma, our approach highlights several signature oncogenes such as PTEN, MYC, CDKN2A, EGFR and FGFR2, which are implicated in controlling cell cycle and angiogenesis (Fig. 5) [70]. In examining how feature attribution shifts when conditioning on morphological features, several genes become more pronounced in predicting survival across each subtype such as ANO9 and RB1 (Appendix C, Fig. 11). ANO9 encodes for a protein that mediates diverse physiological functions such as ion transport and phospholipid movement across the membrane. Over-expression of ANO proteins were found to be correlated with poor prognosis in many tumors, which we similarly observe in our IDHmut ATC subtype with decreased ANO9 expression decreases risk [71]. In addition, we also observe RB1 over-expression decreases risk also in IDHmut ATC, which corroborates with known literature that RB1 is a tumor suppressor gene. Interestingly, EGFR amplification decreased in importance in IDHwt ATC, which may support evidence that EGFR is not a strong therapeutic target in glioblastoma treatment [72]. In CCRCC, Pathomic Fusion discovers decreased CYP3A7 expression and increased PITX2, DDX43, and XIST expression to correlate with risk, which have been linked to cancer predisposition and tumor progression across many cancers (Fig. 6) [73]–[77]. In conditioning on morphological features, HAGHL, MMP1 and ARRP21 gene expression becomes more highly attributed in risk prediction (Appendix C, Fig 11) [78], [79]. For cancers such as CCRCC that do not have multiple molecular subtypes, Pathomic Fusion has the potential to refine gene signatures in cancers, and uncover new prognostic biomarkers that can be targeted in therapeutic treatments.

Across all histology images and cell graphs in both organ types, we observe that Pathomic Fusion broadly localizes both vasculature and cell atypia as an important feature in survival outcome prediction. In ATCs, Pathomic Fusion is able to localize not only regions of tumor cellularity and microvascular proliferation in the histology image, but also glial cells between the microvasculature as depicted in the cell graph (Fig. 5). In ODG, both modalities attend towards “fried egg cells”, which are mildly enlarged round cells with dark nuclei and clear cytoplasm characteristic in ODG. In CCRCC, Pathomic Fusion attends towards cells with indiscernible nucleoli in longer surviving patients, and large cells with clear nucleoli

in shorter surviving patients that is indicative of CCRCC malignancy (Fig. 6). An important aspect about our method is that we are able to leverage heatmaps from both histology images and cell graphs to explain prognostic histological features used for prediction. Though visual explanations from the image and cell graph heatmap often overlap in localizing cell atypia, the cell graph can be used to uncover salient regions that are not recognized in the histology image for risk prediction. Moreover, cell graphs may have additional clinical potential in explainability, as the attributions refer to specific atypical cells rather than pixel regions.

VI. CONCLUSION

The recent advancements made in imaging and sequencing technologies is transforming our understanding of molecular biology and medicine with multimodal data. Next-generation sequencing technologies such as RNA-Seq is redefining clinical grading paradigms to include bulk quantitative measurements from molecular subtyping [59]. Tangential to this growth field has been the emergence of tissue imaging instrumentation such as whole-slide imaging, which capture the organization of cells and their surrounding tissue architecture. In this work, we present Pathomic Fusion, a novel framework for integrating data from these technologies for building objective image-omic assays for cancer diagnosis and prognosis. We extract morphological features from histology images using CNNs and GCNs and genomic features using SNNs and fuse these deep features using the Kronecker Product and a gating-based attention mechanism. We validate our approach on glioma and CCRCC data from TCGA, and demonstrate how multimodal networks in medicine can be used for fine-grained patient stratification and interpreted for finding prognostic features. The method presented is scalable and interpretable for multiple modalities of different data types, and may be used for integrating any combination of imaging and multi-omic data. The paradigm is general and may be used for predicting response and resistance to treatment. Multimodal interpretability has the ability to identify new and novel integrative bio-markers of diagnostic, prognostic and therapeutic relevance.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by the Nvidia GPU grant program and Google Cloud. R.J.C. is funded by the National Science Foundation Graduate Fellowship and the NIH NHGRI T32 training program.

REFERENCES

- [1]. Wen PY and Kesari S, "Malignant gliomas in adults," *New England Journal of Medicine*, vol. 359, no. 5, pp. 492–507, Jul. 2008. [PubMed: 18669428]
- [2]. Nayak L, Lee EQ, and Wen PY, "Epidemiology of brain metastases," *Current Oncology Reports*, vol. 14, no. 1, pp. 48–54, Oct. 2011.
- [3]. Aldape K, Zadeh G, Mansouri S, Reifenberger G, and von Deimling A, "Glioblastoma: pathology, molecular mechanisms and markers," *Acta Neuropathologica*, vol. 129, no. 6, pp. 829–848, May 2015. [PubMed: 25943888]
- [4]. Olar A and Aldape KD, "Using the molecular classification of glioblastoma to inform personalized treatment," *The Journal of Pathology*, vol. 232, no. 2, pp. 165–177, Dec. 2013.

- [5]. Shanes J, Ghali J, Billingham M, Ferrans V, Fenoglio J, Edwards W, Tsai C, Saffitz J, Isner J, and Furner S, "Interobserver variability in the pathologic interpretation of endomyocardial biopsy results." *Circulation*, vol. 75, no. 2, pp. 401–405, 1987. [PubMed: 3802444]
- [6]. Courtiol P, Maussion C, Moarii M, Pronier E, Pilcer S, Sefta M, and et al. , "Deep learning-based classification of mesothelioma improves prediction of patient outcome," *Nature medicine*, vol. 25, no. 10, pp. 1519–1525, 2019.
- [7]. Puchalski RB, Shah N, Miller J, Dalley R, Nomura SR, Yoon J-G, Smith KA, Lankarovich M, Bertagnolli D, Bickley K et al. , "An anatomic transcriptional atlas of human glioblastoma," *Science*, vol. 360, no. 6389, pp. 660–663, 2018. [PubMed: 29748285]
- [8]. Jackson HW, Fischer JR, Zanotelli VR, Ali HR, Mechera R, Soysal SD, Moch H, Muenst S, Varga Z, Weber WP et al. , "The single-cell pathology landscape of breast cancer," *Nature*, vol. 578, no. 7796, pp. 615–620, 2020. [PubMed: 31959985]
- [9]. Schapiro D, Jackson HW, Raghuraman S, Fischer JR, Zanotelli VR, Schulz D, Giesen C, Catena R, Varga Z, and Bodenmiller B, "histocat: analysis of cell phenotypes and interactions in multiplex image cytometry data," *Nature methods*, vol. 14, no. 9, p. 873, 2017. [PubMed: 28783155]
- [10]. Somarakis A, Van Unen V, Koning F, Lelieveldt BP, and Höllt T, "Imacyte: Visual exploration of cellular microenvironments for imaging mass cytometry data," *IEEE transactions on visualization and computer graphics*, 2019.
- [11]. Abdelmoula WM, Balluff B, Englert S, Dijkstra J, Reinders MJ, Walch A, McDonnell LA, and Lelieveldt BP, "Data-driven identification of prognostic tumor subpopulations using spatially mapped t-sne of mass spectrometry imaging data," *Proceedings of the National Academy of Sciences*, vol. 113, no. 43, pp. 12244–12249, 2016.
- [12]. Abdelmoula WM, "Data analysis for mass spectrometry imaging: methods and applications," 2017.
- [13]. Gallego O, "Nonsurgical treatment of recurrent glioblastoma," *Current Oncology*, vol. 22, no. 4, p. 273, May 2015.
- [14]. Yener B, "Cell-graphs: image-driven modeling of structure-function relationship," *Communications of the ACM*, vol. 60, no. 1, pp. 74–84, 2016.
- [15]. Gadiya S, Anand D, and Sethi A, "Histograms: Graphs in histopathology," *arXiv preprint arXiv:1908.05020*, 2019.
- [16]. Wang J, Chen RJ, Lu MY, Baras A, and Mahmood F, "Weakly supervised prostate tma classification via graph convolutional networks," *arXiv preprint arXiv:1910.13328*, 2019.
- [17]. Zuo S, Zhang X, and Wang L, "A RNA sequencing-based six-gene signature for survival prediction in patients with glioblastoma," *Scientific Reports*, vol. 9, no. 1, Feb. 2019.
- [18]. TCGA et al. , "Integrated genomic analyses of ovarian carcinoma," *Nature*, vol. 474, no. 7353, p. 609, 2011. [PubMed: 21720365]
- [19]. Yousefi S, Amrollahi F, Amgad M, Dong C, Lewis JE, Song C, and et al. , "Predicting clinical outcomes from large scale cancer genomic profiles with deep survival models," *Scientific Reports*, vol. 7, no. 1, Sep. 2017.
- [20]. Huang Z, Zhan X, Xiang S, Johnson TS, Helm B, Yu CY, and et al. , "Salmon: Survival analysis learning with multi-omics neural networks on breast cancer," *Frontiers in genetics*, vol. 10, p. 166, 2019. [PubMed: 30906311]
- [21]. Ngiam J, Khosla A, Kim M, Nam J, Lee H, and Ng AY, "Multimodal deep learning," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 689–696.
- [22]. Suresh H, Hunt N, Johnson A, Celi LA, Szolovits P, and Ghassemi M, "Clinical intervention prediction and understanding with deep neural networks," in *Machine Learning for Healthcare Conference*, 2017, pp. 322–337.
- [23]. Kim J-H, On K-W, Lim W, Kim J, Ha J-W, and Zhang B-T, "Hadamard product for low-rank bilinear pooling," *arXiv preprint arXiv:1610.04325*, 2016.
- [24]. Zadeh A, Chen M, Poria S, Cambria E, and Morency L-P, "Tensor fusion network for multimodal sentiment analysis," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2017.

- [25]. Subramanian V, Chidester B, Ma J, and Do MN, "Correlating cellular features with gene expression using cca," in 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). IEEE, 2018, pp. 805–808.
- [26]. Carmichael I and Marron J, "Joint and individual analysis of breast cancer histologic images and genomic covariates," arXiv preprint arXiv:1912.00434, 2019.
- [27]. Coudray N, Ocampo PS, Sakellaropoulos T, Narula N, Snuderl M, Fenyö, and et al. . "Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning," *Nature medicine*, vol. 24, no. 10, p. 1559, 2018.
- [28]. Kather JN, Pearson AT, Halama N, Jäger D, Krause J, Loosen SH, and et al. , "Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer," *Nature medicine*, p. 1, 2019.
- [29]. Mobadersany P, Yousefi S, Amgad M, Gutman DA, Barnholtz-Sloan JS, Vega JEV, and et al. , "Predicting cancer outcomes from histology and genomics using convolutional networks," *Proceedings of the National Academy of Sciences*, vol. 115, no. 13, pp. E2970–E2979, Mar. 2018.
- [30]. Baltrusaitis T, Ahuja C, and Morency L-P, "Multimodal machine learning: A survey and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 2, pp. 423–443, Feb. 2019. [PubMed: 29994351]
- [31]. Cheerla A and Gevaert O, "Deep learning with multimodal representation for pancancer prognosis prediction," *Bioinformatics*, vol. 35, no. 14, pp. i446–i454, Jul. 2019. [PubMed: 31510656]
- [32]. Shao W, Han Z, Cheng J, Cheng L, Wang T, Sun L, Lu Z, Zhang J, Zhang D, and Huang K, "Integrative analysis of pathological images and multi-dimensional genomic data for early-stage cancer prognosis," *IEEE Transactions on Medical Imaging*, 2019.
- [33]. Prewitt JM, "Graphs and grammars for histology: An introduction," in *Proceedings of the Annual Symposium on Computer Application in Medical Care*. American Medical Informatics Association, 1979, p. 18.
- [34]. Doyle S, Hwang M, Shah K, Madabhushi A, Feldman M, and Tomaszewski J, "Automated grading of prostate cancer using architectural and textural image features," in 2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro. IEEE, 2007, pp. 1284–1287.
- [35]. Hamilton W, Ying Z, and Leskovec J, "Inductive representation learning on large graphs," in *Advances in Neural Information Processing Systems*, 2017, pp. 1024–1034.
- [36]. Lee J, Lee I, and Kang J, "Self-attention graph pooling," arXiv preprint arXiv:1904.08082, 2019.
- [37]. Defferrard M, Bresson X, and Vandergheynst P, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in neural information processing systems*, 2016, pp. 3844–3852.
- [38]. Kipf TN and Welling M, "Semi-supervised classification with graph convolutional networks," arXiv preprint arXiv:1609.02907, 2016.
- [39]. Zhou Y, Graham S, Koohbanani NA, Shaban M, Heng P-A, and Rajpoot N, "Cgc-net: Cell graph convolutional network for grading of colorectal cancer histology images," arXiv preprint arXiv:1909.01068, 2019.
- [40]. Marusyk A, Almendro V, and Polyak K, "Intra-tumour heterogeneity: a looking glass for cancer?" *Nature Reviews Cancer*, vol. 12, no. 5, pp. 323–334, Apr. 2012. [PubMed: 22513401]
- [41]. Yang H, Kim J-Y, Kim H, and Adhikari SP, "Guided soft attention network for classification of breast cancer histopathology images," *IEEE transactions on medical imaging*, 2019.
- [42]. Sari CT and Gunduz-Demir C, "Unsupervised feature extraction via deep learning for histopathological classification of colon tissue images," *IEEE transactions on medical imaging*, vol. 38, no. 5, pp. 1139–1149, 2018. [PubMed: 30403624]
- [43]. Bandi P, Geessink O, Manson Q, Van Dijk M, Balkenhol M, Hermsen M, and et al. , "From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge," *IEEE transactions on medical imaging*, vol. 38, no. 2, pp. 550–560, 2018.

- [44]. Guadagno E, Borrelli G, Califano M, Cali G, Solari D, and Caro MDBD, "Immunohistochemical expression of stem cell markers CD44 and nestin in glioblastomas: Evaluation of their prognostic significance," *Pathology - Research and Practice*, vol. 212, no. 9, pp. 825–832, Sep. 2016. [PubMed: 27450656]
- [45]. Mikkelsen VE, Stensjøen AL, Berntsen EM, Nordrum IS, Salvesen Ø, Solheim O, and Torp SH, "Histopathologic features in relation to pretreatment tumor growth in patients with glioblastoma," *World Neurosurgery*, vol. 109, pp. e50–e58, Jan. 2018. [PubMed: 28951271]
- [46]. Naylor P, Laé M, Reyat F, and Walter T, "Segmentation of nuclei in histopathology images by deep regression of the distance map," *IEEE transactions on medical imaging*, vol. 38, no. 2, pp. 448–459, 2018.
- [47]. Xu J, Xiang L, Liu Q, Gilmore H, Wu J, Tang J, and Madabhushi A, "Stacked sparse autoencoder (ssae) for nuclei detection on breast cancer histopathology images," *IEEE transactions on medical imaging*, vol. 35, no. 1, pp. 119–130, 2015. [PubMed: 26208307]
- [48]. Su H, Xing F, and Yang L, "Robust cell detection of histopathological brain tumor images using sparse reconstruction and adaptive dictionary selection," *IEEE transactions on medical imaging*, vol. 35, no. 6, pp. 1575–1586, 2016. [PubMed: 26812706]
- [49]. Jia Z, Huang X, Eric I, Chang C, and Xu Y, "Constrained deep weak supervision for histopathology image segmentation," *IEEE transactions on medical imaging*, vol. 36, no. 11, pp. 2376–2388, 2017. [PubMed: 28692971]
- [50]. Kumar N, Verma R, Sharma S, Bhargava S, Vahadane A, and Sethi A, "A dataset and a technique for generalized nuclear segmentation for computational pathology," *IEEE Transactions on Medical Imaging*, vol. 36, no. 7, pp. 1550–1560, Jul. 2017. [PubMed: 28287963]
- [51]. Mahmood F, Borders D, Chen R, McKay G, Salimian KJ, Baras A, , and Durr NJ, "Adversarial training for multi-organ nuclei segmentation in computational pathology images," in *IEEE Transactions on Medical Imaging*, 2018.
- [52]. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, and Bengio Y, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [53]. Isola P, Zhu J-Y, Zhou T, and Efros AA, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [54]. Muja M and Lowe DG, "Fast approximate nearest neighbors with automatic algorithm configuration." *VISAPP (1)*, vol. 2, no. 331-340, p. 2, 2009.
- [55]. Oord A. v. d., Li Y, and Vinyals O, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [56]. Hénaff OJ, Razavi A, Doersch C, Eslami S, and Oord A. v. d., "Data-efficient image recognition with contrastive predictive coding," *arXiv preprint arXiv:1905.09272*, 2019.
- [57]. Lu MY, Chen RJ, Wang J, Dillon D, and Mahmood F, "Semi-supervised histology classification using deep multiple instance learning and contrastive predictive coding," *arXiv preprint arXiv:1910.10825*, 2019.
- [58]. Klambauer G, Unterthiner T, Mayr A, and Hochreiter S, "Self-normalizing neural networks," in *Advances in neural information processing systems*, 2017, pp. 971–980.
- [59]. Louis DN, Perry A, Reifenberger G, von Deimling A, Figarella-Branger D, Cavenee WK, and et al. , "The 2016 world health organization classification of tumors of the central nervous system: a summary," *Acta Neuropathologica*, vol. 131, no. 6, pp. 803–820, May 2016. [PubMed: 27157931]
- [60]. Arevalo J, Solorio T, Montes M, and Gonzalez F, "Gated multimodal units for information fusion," 02 2017.
- [61]. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, and Batra D, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [62]. Sundararajan M, Taly A, and Yan Q, "Axiomatic attribution for deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70. JMLR. org*, 2017, pp. 3319–3328.

- [63]. Bach S, Binder A, Montavon G, Klauschen F, Müller K-R, and Samek W, “On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation,” *PloS one*, vol. 10, no. 7, p. e0130140, 2015. [PubMed: 26161953]
- [64]. Vahadane A, Peng T, Sethi A, Albarqouni S, Wang L, Baust M, and et al. , “Structure-preserving color normalization and sparse stain separation for histological images,” *IEEE Transactions on Medical Imaging*, vol. 35, no. 8, pp. 1962–1971, Aug. 2016. [PubMed: 27164577]
- [65]. Cerami E, , Gao J, Dogrusoz U, Gross BE, Sumer SO, Aksoy BA, and et al. , “The cBio cancer genomics portal: An open platform for exploring multidimensional cancer genomics data: Figure 1.” *Cancer Discovery*, vol. 2, no. 5, pp. 401–404, May 2012. [PubMed: 22588877]
- [66]. Shuch B, Hofmann JN, Merino MJ, Nix JW, Vourganti S, Linehan WM, Schwartz K, Ruterbusch JJ, Colt JS, Purdue MP et al. , “Pathologic validation of renal cell carcinoma histology in the surveillance, epidemiology, and end results program,” in *Urologic Oncology: Seminars and Original Investigations*, vol. 32, no. 1. Elsevier, 2014, pp. 23–e9.
- [67]. Linehan WM, Walther MM, and Zbar B, “The genetic basis of cancer of the kidney,” *The Journal of urology*, vol. 170, no. 6, pp. 2163–2172, 2003. [PubMed: 14634372]
- [68]. Schraml P, Struckmann K, Hatz F, Sonnet S, Kully C, Gasser T, Sauter G, Mihatsch MJ, and Moch H, “Vhl mutations and their correlation with tumour cell proliferation, microvessel density, and patient prognosis in clear cell renal cell carcinoma,” *The Journal of Pathology: A Journal of the Pathological Society of Great Britain and Ireland*, vol. 196, no. 2, pp. 186–193, 2002.
- [69]. Bland JM and Altman DG, “The logrank test,” *Bmj*, vol. 328, no. 7447, p. 1073, 2004. [PubMed: 15117797]
- [70]. “Comprehensive, integrative genomic analysis of diffuse lower-grade gliomas,” *New England Journal of Medicine*, vol. 372, no. 26, pp. 2481–2498, Jun. 2015. [PubMed: 26061751]
- [71]. Jun I, Park HS, Piao H, Han JW, An MJ, Yun BG, Zhang X, Cha YH, Shin YK, Yook JI et al. , “Ano9/tmem16j promotes tumorigenesis via egfr and is a novel therapeutic target for pancreatic cancer,” *British journal of cancer*, vol. 117, no. 12, pp. 1798–1809, 2017. [PubMed: 29024940]
- [72]. Westphal M, Maire CL, and Lamszus K, “EGFR as a target for glioblastoma treatment: An unfulfilled promise,” *CNS Drugs*, vol. 31, no. 9, pp. 723–735, Aug. 2017. [PubMed: 28791656]
- [73]. Murray G, McFadyen M, Mitchell R, Cheung Y, Kerr A, and Melvin W, “Cytochrome p450 cyp3a in human renal cell cancer,” *British journal of cancer*, vol. 79, no. 11, pp. 1836–1842, 1999. [PubMed: 10206301]
- [74]. Abdel-Fatah T, McArdle S, Johnson C, Moseley P, Ball G, Pockley A, Ellis I, Rees R, and Chan S, “Hage (ddx43) is a biomarker for poor prognosis and a predictor of chemotherapy response in breast cancer,” *British journal of cancer*, vol. 110, no. 10, pp. 2450–2461, 2014. [PubMed: 24755885]
- [75]. Allegra A, Alonci A, Campo S, Penna G, Petrunaro A, Gerace D, and Musolino C, “Circulating micrnas: new biomarkers in diagnosis, prognosis and treatment of cancer,” *International journal of oncology*, vol. 41, no. 6, pp. 1897–1912, 2012. [PubMed: 23026890]
- [76]. Jordan AR, Wang J, Yates TJ, Hasanali SL, Lokeshwar SD, Morera DS, Shamaladevi N, Li CS, Klaassen Z, Terris MK et al. , “Molecular targeting of renal cell carcinoma by an oral combination,” *Oncogenesis*, vol. 9, no. 5, pp. 1–13, 2020. [PubMed: 31900385]
- [77]. Lee W-K and Thévenod F, “Oncogenic pitx2 facilitates tumor cell drug resistance by inverse regulation of hoct3/slc22a3 and abc drug transporters in colon and kidney cancers,” *Cancer letters*, vol. 449, pp. 237–251, 2019. [PubMed: 30742940]
- [78]. Lin T-C, Yeh Y-M, Fan W-L, Chang Y-C, Lin W-M, Yang T-Y, and Hsiao M, “Ghrelin upregulates oncogenic aurora a to promote renal cell carcinoma invasion,” *Cancers*, vol. 11, no. 3, p. 303, 2019. [PubMed: 30836712]
- [79]. Wang Q-M, Lv L, Tang Y, Zhang L, and Wang L-F, “Mmp-1 is overexpressed in triple-negative breast cancer tissues and the knockdown of mmp-1 expression inhibits tumor cell malignant behaviors in vitro,” *Oncology letters*, vol. 17, no. 2, pp. 1732–1740, 2019. [PubMed: 30675232]

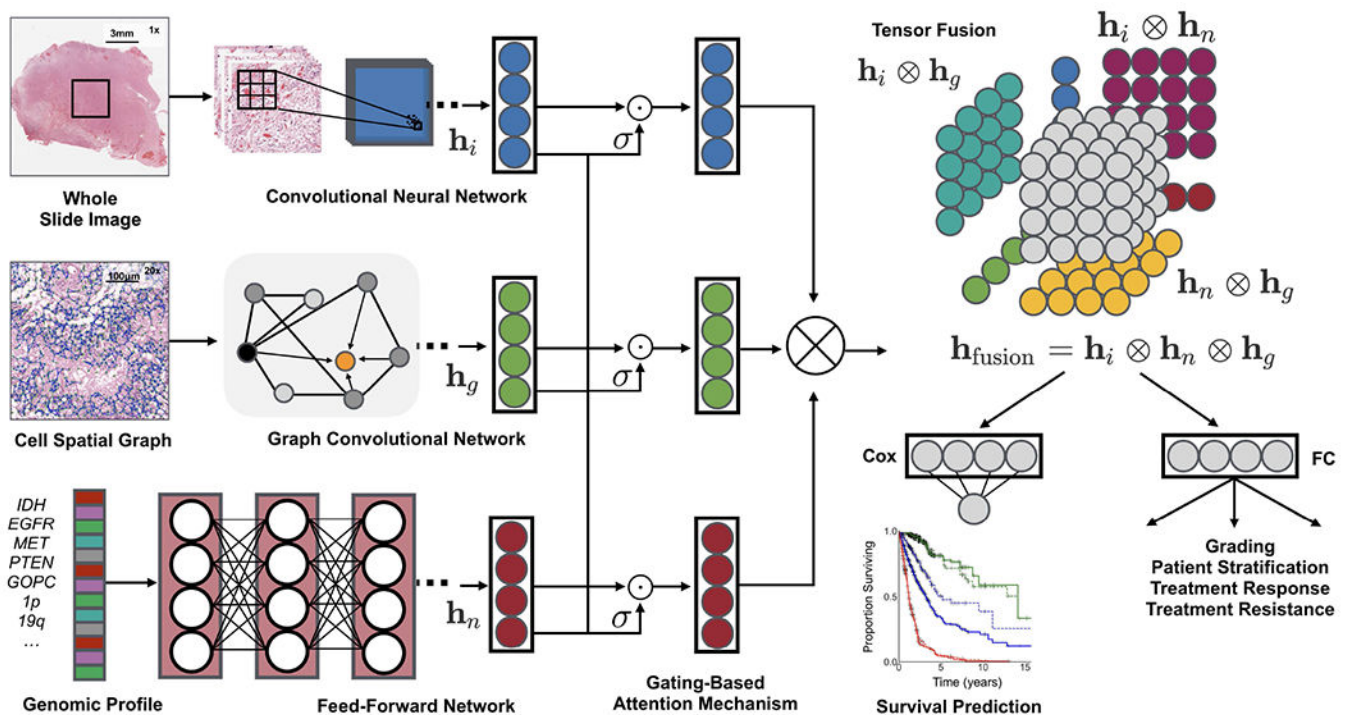


Fig. 1: Pathomic Fusion:

An integrated framework for multimodal fusion of histology and genomic features for survival outcome prediction and classification. Histology features may be extracted using CNNs, parameter efficient GCNs or a combination of the two. Unimodal networks for the respective image and genomic features are first trained individually for the corresponding supervised learning task, and then used as feature extractors for multimodal fusion. Multimodal fusion is performed by applying an gating-based attention mechanism to first control the expressiveness of each modality, followed by the Kronecker product to model pairwise feature interactions across modalities.

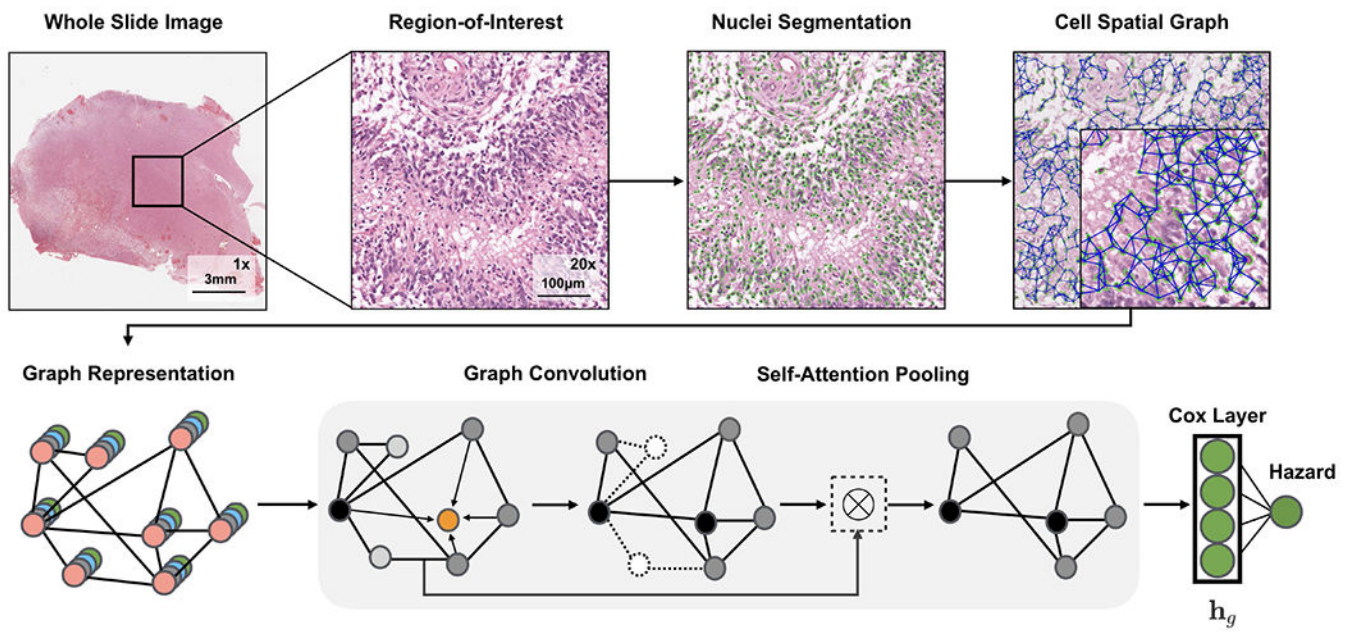


Fig. 2: Graph Convolutional Network for learning morphometric cell features from histology images. We represent cells in histology tissue as nodes in a graph, where cells are isolated using a deep learning-based nuclei segmentation algorithm and the connections between cells are made using KNN. Features for each cell are initialized using handcrafted features as well as deep features learned using contrastive predictive coding. The aggregate and combine functions are adopted from the GraphSAGE architecture, with the node masking and hierarchical pooling strategy adopted from SAGEPool.

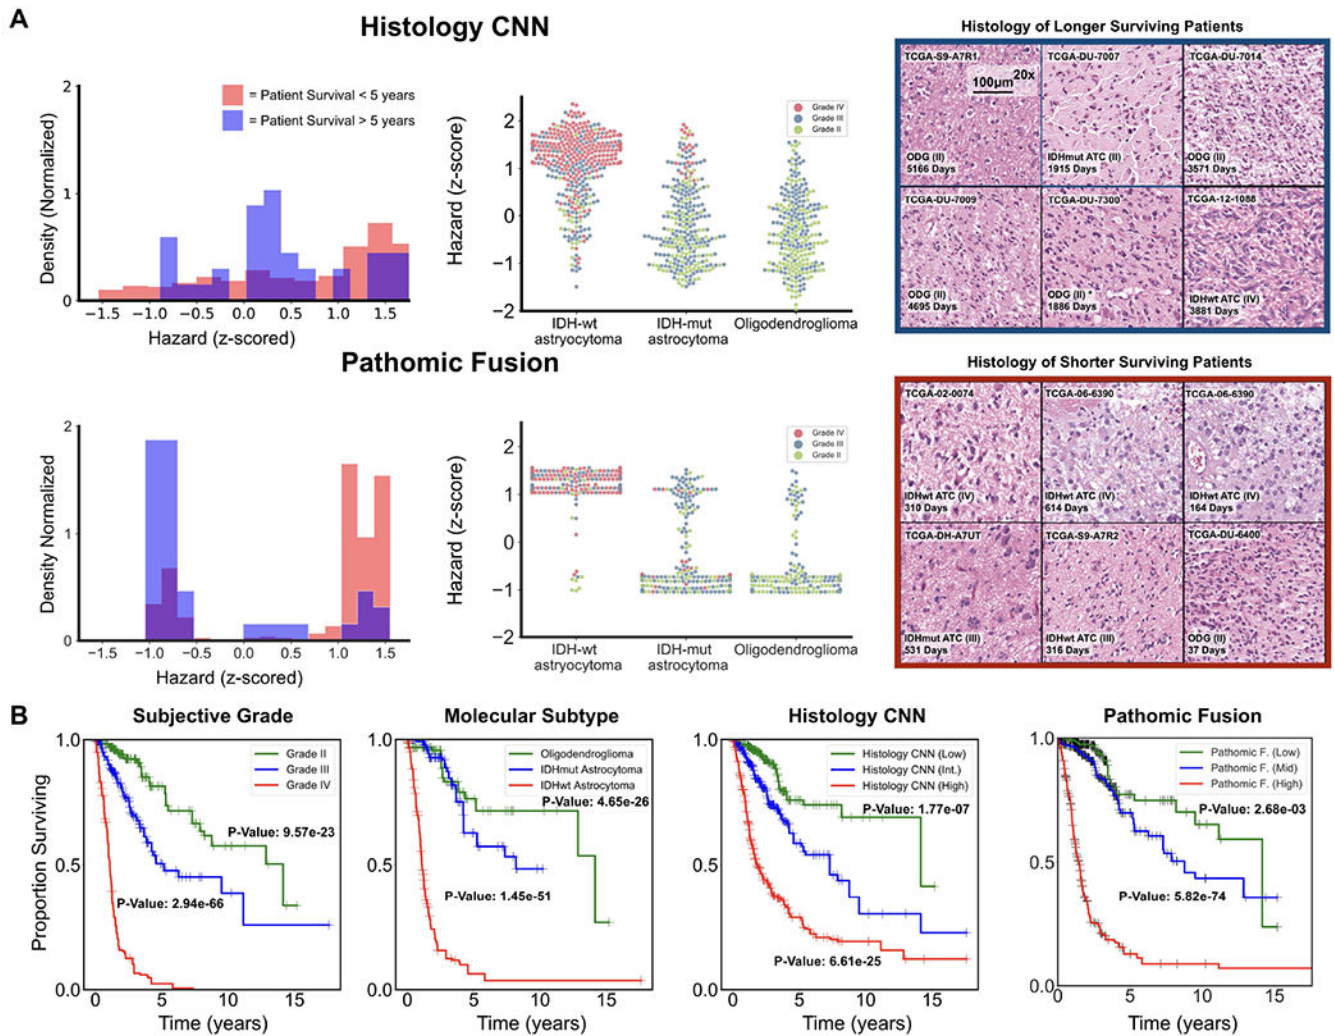


Fig. 3: Pathomic Fusion Applied to Glioblastoma and Lower Grade Glioma.

A. Glioma hazard distributions amongst shorter vs. longer surviving uncensored patients and molecular subtypes for Histology CNN and Pathomic Fusion. Patients are defined as shorter surviving if patient death is observed before 5 years of the first follow-up (shaded red), and longer surviving if patient death is observed after 5 years of the first follow-up (shaded blue). Pathomic Fusion predicts hazard in more concentrated clusters than Histology CNN, while the distribution of hazard predictions from Histology CNN have longer tails and are more varied across molecular subtypes. In analyzing the types of glioma in the three high density regions revealed from Pathomic Fusion, we see that these regions corroborate with the WHO paradigm for stratifying patients into IDHwt ATC, IDHmut ATC, and ODG (Appendix C, Table IV). **B.** Kaplan-Meier comparative analysis of using grade, molecular subtype, Histology CNN and Pathomic Fusion in stratifying patient outcomes. Hazard predictions from Pathomic Fusion show better stratification of mid-to-high risk patients than Histology CNN, and low-to-mid risk patients than molecular subtyping, which follows the WHO paradigm. Low / intermediate / high risk are defined by the 33-66-100 percentile

of hazard predictions. Overlaid Kaplan-Meier estimates of our network predictions with WHO Grading is shown in the supplement (Appendix C, Fig. 9).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

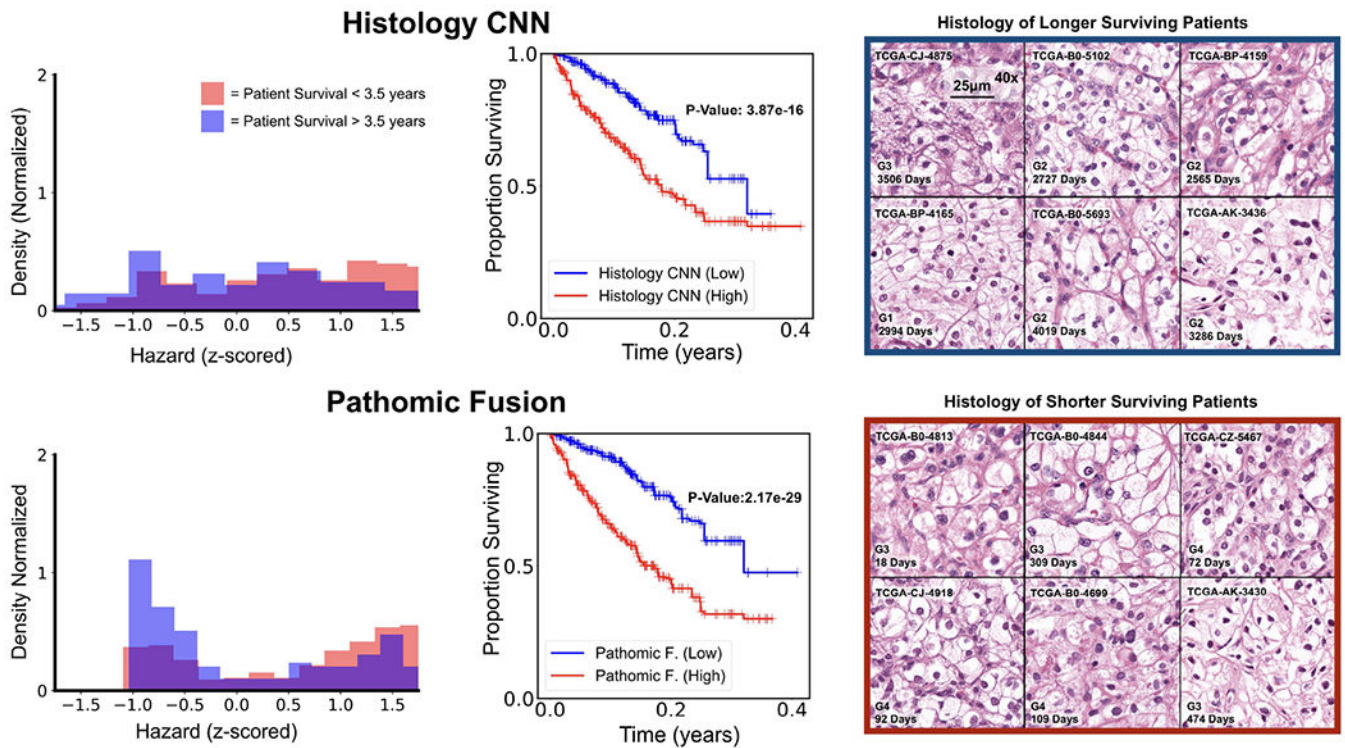


Fig. 4: Pathomic Fusion Applied to Clear Cell Renal Cell Carcinoma.

CCRCC hazard distributions amongst shorter vs. longer surviving uncensored patients for Histology CNN and Pathomic Fusion. Patients are defined as shorter surviving if patient death is observed before 3.5 years of the first follow-up (shaded red), and longer surviving if patient death is observed after 3.5 years of the first follow-up (shaded blue). Pathomic Fusion was observed to be able to stratify longer and shorter surviving patients better than Histology CNN, exhibiting a bimodal distribution in hazard prediction. Overlaid Kaplan-Meier estimates of our network predictions with WHO Grading is shown in the supplement (Appendix C, Fig. 10).

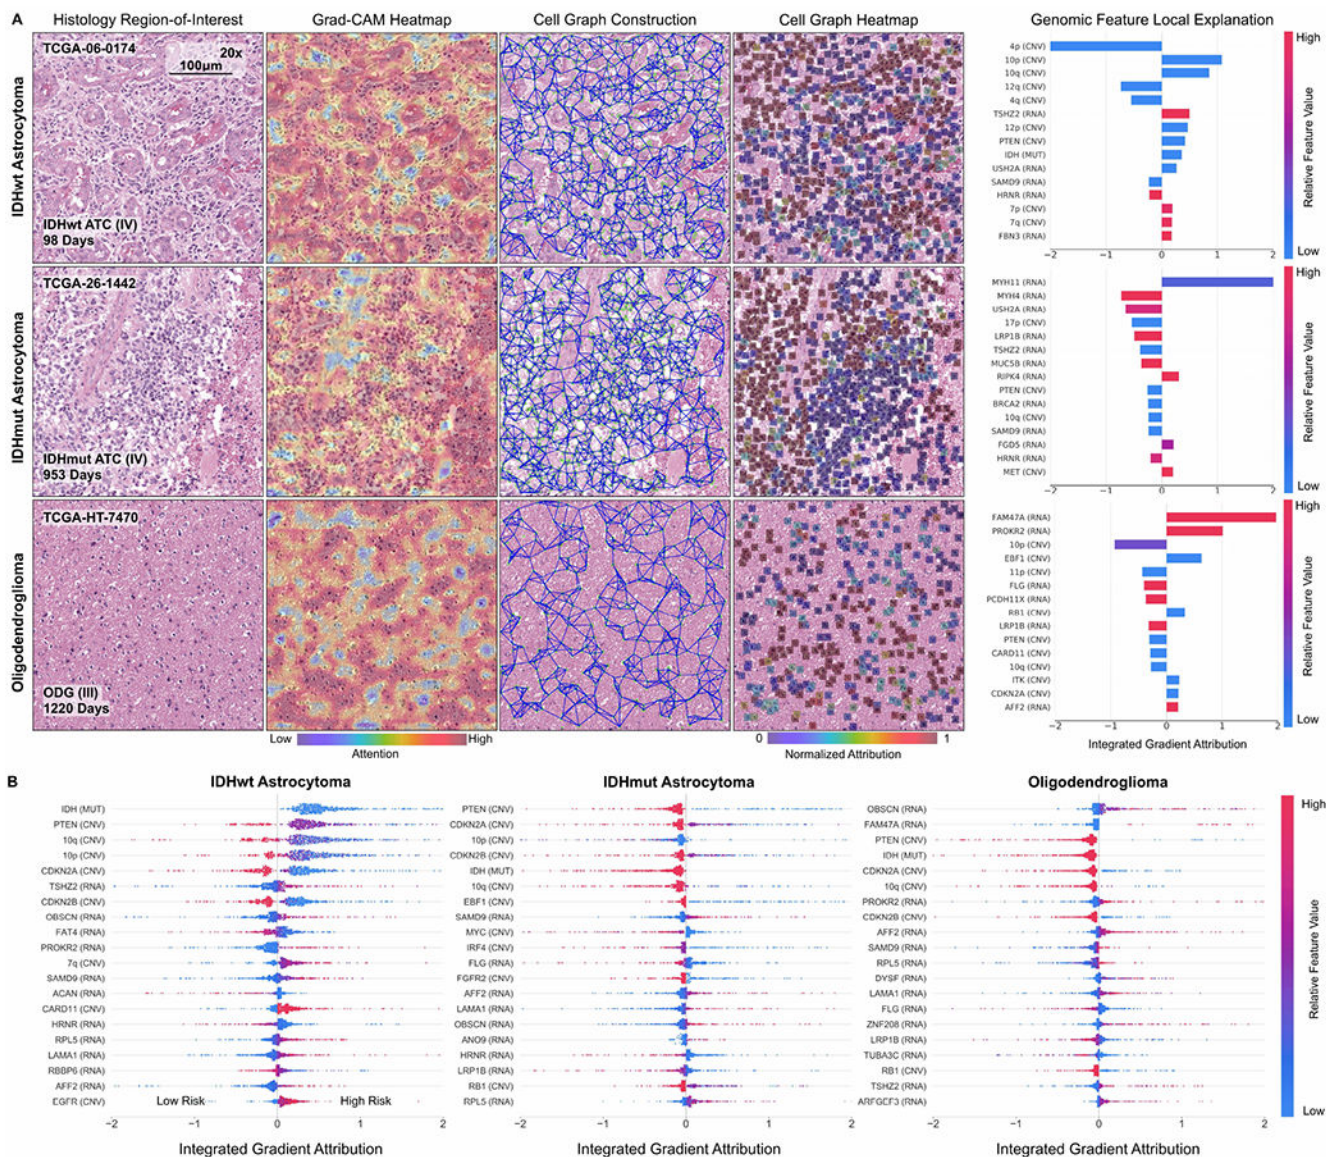


Fig. 5: Multimodal interpretability by Pathomic Fusion in glioma. **A.** Local explanation of histology image, cell graph, and genomic modalities for individual patients of three molecular subtypes. In IDHwt ATC, the network detects endothelial cells of the microvascular proliferation in the histology image, while the cell graph localizes glial cells between the microvasculature. In IDHmut ATC, we observe similar localization of tumor cellularity in both the histology image and cell graph, however, attribution direction for IDH is flipped to have positive impact on survival. In ODG, we observe both modalities localizing towards different regions containing “fried egg cells” that are canonical in ODG. For each of these patients, local explanation reveals the most important genomic features used for prediction. **B.** Global explanation of top 20 genomic features for each molecular subtype in glioma. Canonical oncogenes in glioma such as IDH, PTEN, MYC and CDKN2A are attributed highly as being important for risk prediction.

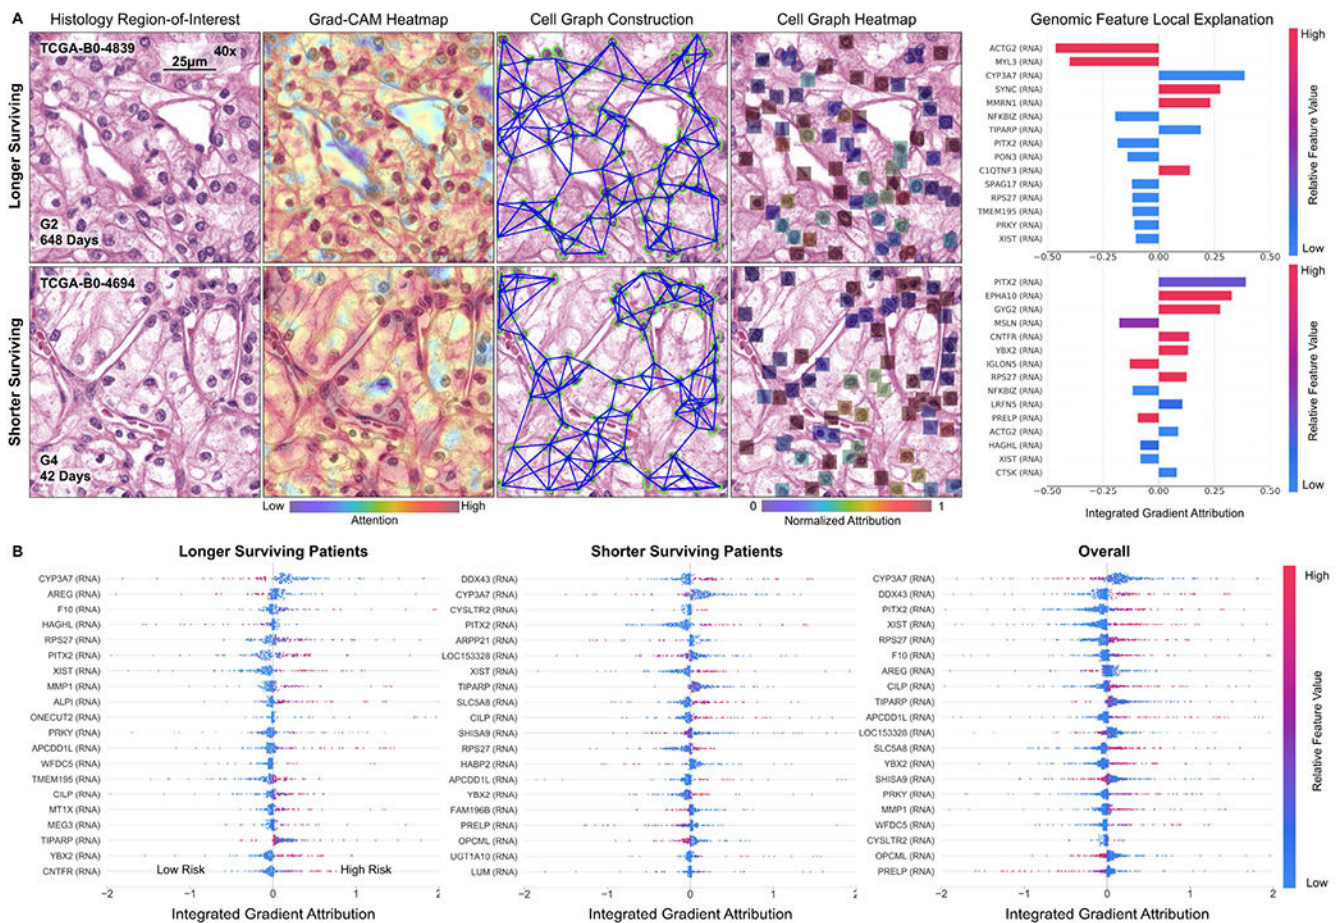


Fig. 6: Multimodal interpretability by Pathomic Fusion in CCRCC. **A.** Local explanation of histology image, cell graph, and genomic modalities for two longer and shorter surviving patients. In the longer surviving patient, Pathomic Fusion localizes cells without obvious nucleoli in both the histology image and cell graph, which suggests lower-grade CCRCC and lower risk. In the shorter surviving patient, we observe Pathomic Fusion attending to large cells with prominent nucleoli and eosinophilic-to-clear cytoplasm in the cell graph, and the “chicken-wire” vasculature pattern in the histology image that is characteristic of higher-grade CCRCC. Cells without clear cytoplasm are noticeably missed in both modalities for shorter survival. For each of these patients, local explanation reveals the most important genomic features used for prediction. **B.** Global explanation of top 20 genomic features for longer surviving, shorter surviving, and all patients in CCRCC. Genes such as CYP3A7, DDX43 and PITX2 are attributed highly as being important for risk prediction, which have linked to cancer predisposition and tumor progression in CCRCC and other cancers.

TABLE I:

Concordance Index of Pathomic Fusion and ablation experiments in glioma survival prediction.

Model	c-Index
Cox (Age+Gender)	0.732 ± 0.012 *
Cox (Grade)	0.738 ± 0.013 *
Cox (Molecular Subtype)	0.760 ± 0.011 *
Cox (Grade+Molecular Subtype)	0.777 ± 0.013 *
Histology CNN	0.792 ± 0.014 *
Histology GCN	0.746 ± 0.023 *
Genomic SNN	0.808 ± 0.014 *
SCNN (Histology Only) [29]	0.754 *
GSCNN (Histology + Genomic) [29]	0.781 *
<i>Pathomic F.</i> (GCN⊗SNN)	0.812 ± 0.010 *
<i>Pathomic F.</i> (CNN⊗SNN)	0.820 ± 0.009 *
<i>Pathomic F.</i> (CNN⊗GCN⊗SNN)	0.826 ± 0.009 *

*
 $p < 0.05$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE II:

Concordance Index of Pathomic Fusion and ablation experiments in CCRCC survival prediction.

Model	c-Index
Cox (Age+Gender)	0.630 ± 0.024 *
Cox (Grade)	0.675 ± 0.036 *
Histology CNN	0.671 ± 0.023 *
Histology GCN	0.646 ± 0.022 *
Genomic SNN	0.684 ± 0.025 *
<i>Pathomic F. (GCN⊗SNN)</i>	0.688 ± 0.029 *
<i>Pathomic F. (CNN⊗SNN)</i>	0.719 ± 0.031 *
<i>Pathomic F. (CNN⊗GCN⊗SNN)</i>	0.720 ± 0.028 *

*
 $p < 0.05$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript