

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# Multi stain graph fusion for multimodal integration in pathology

Chaitanya Dwivedi<sup>1</sup>, Shima Nofallah<sup>1</sup>, Maryam Pouryahya<sup>1</sup>, Janani Iyer<sup>1</sup>, Kenneth Leidal<sup>1</sup>, Chuhan Chung<sup>2</sup>, Timothy Watkins<sup>2</sup>, Andrew Billin<sup>2</sup>, Robert Myers<sup>2</sup>, John Abel<sup>1</sup>, and Ali Behrooz<sup>1</sup>

> <sup>1</sup>PathAI, Inc., <sup>2</sup>Gilead Sciences, Inc. {chaitanya.dwivedi, ali.behrooz}@pathai.com

### Abstract

In pathology, tissue samples are assessed using multiple staining techniques to enhance contrast in unique histologic features. In this paper, we introduce a multimodal CNN-GNN based graph fusion approach that leverages complementary information from multiple non-registered histopathology images to predict pathologic scores. We demonstrate this approach in nonalcoholic steatohepatitis (NASH) by predicting CRN fibrosis stage and NAFLD Activity Score (NAS). Primary assessment of NASH typically requires liver biopsy evaluation on two histological stains: Trichrome (TC) and hematoxylin and eosin (H&E). Our multimodal approach learns to extract complementary information from TC and H&E graphs corresponding to each stain while simultaneously learning an optimal policy to combine this information. We report up to 20% improvement in predicting fibrosis stage and NAS component grades over single-stain modeling approaches, measured by computing linearly weighted Cohen's kappa between machinederived vs. pathologist consensus scores. Broadly, this paper demonstrates the value of leveraging diverse pathology images for improved ML-powered histologic assessment.

## 1. Introduction

Graph neural networks (GNN) are increasingly used in digital pathology, as they enable the integration of tissue spatial structure into the prediction of clinically relevant metrics [1, 2, 6]. Graph representation of a tissue sample can be built using features extracted from a single digital whole slide image (WSI). However, in many cases information from multiple images may be necessary for disease assessment. These images are differently stained to identify specific molecular features; moreover, they may be taken from different areas of the diseased tissue. Thus, combining information from multiple images is rarely as straightforward as aligning serial sections to make a single graph, and approaches to graph fusion are needed.

One specific use case of such a graph fusion approach

is in the evaluation of tissue biopsies from nonalcoholic steatohepatitis (NASH). NASH, the progressive form of non-alcoholic fatty liver disease, is increasingly prevalent globally, and is characterized by steatosis (fat accumulation), inflammation, and fibrosis leading to liver damage and potentially the need for liver transplant as it grows in severity. Studies show that around 20-50% of the US population suffers from early and late stages of NASH, with risk factors including obesity, diabetes, and metabolic syndrome [48].

Primary medical evaluation of NASH, both clinically and in clinical trials, typically involves histological review of liver biopsies by histopathologists. In NASH histology, tissue samples are assessed using two staining techniques that provide contrast for relevant features: hematoxylin and eosin (H&E) and Masson's trichrome (TC). H&E stains cell nuclei and cytoplasmic features, whereas TC differentiates collagen fibers against nuclei and cytoplasm [4, 50]. NASH disease severity is commonly scored on the ordinal grading system developed by NASH Clinical Research Network (CRN) where H&E- and TC-stained samples are evaluated separately: H&E-stained tissue is assessed for the degree of steatosis, inflammation, and ballooning to generate a composite NAFLD Activity score, and TC-stained samples are evaluated for the level of fibrosis within to render a fibrosis stage of 0 to 4.

While previous works have successfully demonstrated the benefits of a quantitative, reproducible, and automated ML system that analyzes liver biopsy WSIs for the evaluation of NASH to a high degree of concordance with human pathologists [16, 44, 47], they extract information from only one of the two stains routinely available for examining such samples. The contributions of this paper are two-fold:

- We show that combining information from H&E and TC whole-slide images leads to improved model performance for predicting NAS component and fibrosis scores with higher concordance over models trained using single-stain images.
- 2. We propose a global graph attention-based fusion ap-

proach that allows for information to flow between the two graphs during graph convolution. This is in contrast to the common late fusion approaches where information from the two graphs is combined at the end of graph convolution. The global attention inter message passing (GAIMP) architecture outperforms nonfusion approaches.

Importantly, we show these results using diverse, uncurated, and imbalanced data acquired from clinical trials, providing further evidence that our method yields benefit in real-world conditions.

## 2. Related work

Today, there is more demand for AI-based tools for histopathology and digital whole slide image analysis. This is primarily attributed to inter- and intra-rater variability in pathology and widespread adoption of digital scanners in clinics and laboratories [1,27,33].

### 2.1. Deep learning methods for WSI analysis

Deep learning has achieved unparalleled success in histopathology image analysis and enabled the fast, accurate, and robust classification of complex cell and tissue morphologies. A whole slide image (WSI) is produced by scanning a glass slide containing a stained sample of diseased tissue at a microscopic resolution (~0.25-0.5 microns per pixel). The high-resolution scan captures sufficiently detailed microscopic information necessary for pathology disease analysis. The size of WSIs is on the order of  $100,000 \times 100,000$  pixels taking up several hundred megabytes of memory. Pathologists are tasked with meticulously scanning through these large images at several magnification levels to look for relevant biological structures that enable diagnosis, prognosis, or treatment decisions. However, manual assessment in pathology is subject to inter-observer variability which may affect patient treatment or outcomes [37].

Common deep learning methods used for WSI analysis can be broadly binned into two categories:

Segmentation Given the large size of WSIs, convolutional neural networks (CNNs) are applied on smaller patches sampled from WSIs. Patch-level segmentation models make predictions about a single patch without any global context of the larger WSI. Wang *et al.* [46] developed a CNN-based metastatic cancer detection system that achieves an AUC of 0.925 for WSI classification and AUC of 0.7 for tumor localization. Andrew Janowczyk and Anant Madabhushi [21] used AlexNet [28] in nuclei segmentation, tubule segmentation, lymphocyte detection, and lymphoma classification. Taylor-Weiner *et al.* [44] built a CNN-based

WSI segmentation system to measure key histological features in NASH including steatosis, inflammation, hepatocellular ballooning, and fibrosis. The pixel-level features were used to calculate features summarizing the entire segmentation maps. The authors showed strong concordance between scores generated from their system and expert pathologists. Additionally, they evaluated the prognostic capabilities of their system for liver-related clinical events. Heinemann *et al.* [16] used Masson's trichrome WSIs to develop 4 CNNs that classify a patch into fibrosis, ballooning, inflammation, and steatosis stages. Slide level classes were obtained by averaging the logits from each model over all the patches of a trichrome WSI.

2. Slide-level prediction Deep learning models can be used to aggregate pixel-level features obtained from segmentation models to make slide-level predictions. Iizuka *et al.* [18] use CNN-RNN architecture to classify entire WSIs into Adenocarcinoma, Adenoma, and Non-neoplastic. Ilse *et al.* [19] employ a CNN-based multiple instance learning framework [10] for WSI classification in colon and breast cancer. Similarly, Gadiya *et al.* [13] classify WSIs into cancerous vs non-cancerous and in situ vs invasive by representing tissue section as a multi-attributed spatial graph of cells. Wang *et al.* [47] construct graph representations of liver biopsy WSIs using pixel-level CNN features to classify WSIs into different stages of NAS and fibrosis.

#### 2.2. Multi modal integration for WSI analysis

In practice, medical prognostic and diagnostic decisions are often made by collating information from multiple sources including various kinds of tests that uniquely characterize the disease, e.g., lab tests, imaging tests (CT Scan, MRI, etc.), and tissue biopsy (needle, surgery, endoscopy) [20]. Information from a variety of modalities provide complementary features that are independently predictive of an endpoint. Deep learning systems have shown to benefit from integration of such multi-modal information.

Mobadersany *et al.* [35] showed that concatenation of latent histology and genomic features leads to models better predictive of survival outcome in glioma. Chen *et al.* [8] developed an interpretable latent Kronecker product-based strategy for end-to-end multimodal fusion of histologic image and genomic features for survival outcome prediction. Chen *et al.* [7] integrated quantitative latent features from WSIs with mutations, transcriptomics, and proteomics data for prognosis prediction in lung adenocarcinoma. They showed improved survival prediction for LUAD using the multimodal approach. Recently, Chen *et al.* [9] developed a transformer-based multimodal integration strategy that captures interactions between histology-based visual concepts

and gene features.

While integration of WSI features with multi-omics data has been extensively studied previously, to the best of our knowledge, this paper is the first to explore the multimodal integration of multiple histology stains.

### 3. Methods

We train CNN models to perform pixel-level classification on WSIs of H&E and TC stains. We use the heatmaps output by CNNs to generate graphs. The GNN models are trained on these sets of graphs. Following is a detailed explanation of the entire modeling pipeline. The process of generating tissue heatmaps and the construction of graphs is summarized in sections 3.1 and 3.2, respectively. In section 3.3 we describe the details of the single graph GNNs used as baselines. Section 3.4 discusses different fusion methods used to combine information from both H&E and TC stains.

#### 3.1. ML model to generate tissue heatmap on WSIs

Taylor-Weiner et al. [44] presented CNN-based segmentation models which generate tissue heatmaps on NASH WSIs. Similar to [44] we train two CNN models, one for H&E slides and another for TC slides, to perform patch classification. Each model learns to classify the center pixel of the patch into key categories of histological morphologies, 13 for H&E and 5 for TC. During inference, similar to Long *et al.* [32], the fully connected classification layer was converted to a convolutional layer. Transforming fully connected layers into convolution layers enables a classification net to output a spatial map. Background and artifactcontaining regions were excluded from the analysis using additional models trained to classify pixels as either background, tissue with artifact, or usable tissue. All other models and features were then evaluated only in the areas classified as usable tissue. The network architecture is described in more detail in section 4.2.

### **3.2. Graph Construction**

Graphs are constructed based on heatmaps generated by the CNN models. We randomly sample pixels from the heatmaps and use the pixel coordinates and the class logits to cluster pixels into 5000 groups. Clustering is done using the Birch algorithm [51]. Each cluster in the heatmap corresponds to one node in the graph and the node features are derived from pixel features within the cluster. We use three types of features computed from segmentation heatmaps. These features are pre-defined to represent biological morphologies with known clinical relevance [47]. The spatial features include the mean and standard deviation of spatial coordinates of all the cluster pixels. Topological features include the area, perimeter, and convexity of the cluster, and logit-related features include the mean and standard deviation of logits for each of the classes (13 H&E and 5 TC) corresponding to all the cluster pixels. For every node, one directed edge is added to connect to each of its five nearest neighbors.

### 3.3. Single Stain GNNs

Graphs generated in the previous section are used to train non-fusion stain-specific GNNs. This section describes the structure of GNNs.

GNNs can be considered a variant of CNNs that operate directly on graphs [25]. During the training process of a GNN, feature vectors on each node iteratively aggregate feature vectors from their neighbor node as a form of message passing to generate a new feature vector at the next hidden layer in the model. The message passing operation is similar to the convolution operation in CNNs and can be defined as follows:

$$a_v^{(k)} = AGGREGATE^{(k)}(\{h_u^{(k-1)} : u \in \mathcal{N}(v)\}) \quad (1)$$

where  $h_u^{(k)}$  is the feature vector of node v at the k-1-th iteration of the neighborhood aggregation and  $a_v^{(k)}$  is the feature vector of node v at the next iteration. The neighborhood  $\mathcal{N}(v) = \{u \in V | (v, u) \in E\}$  of node v is the set of adjacent nodes of v.

Using a pooling operation over all the nodes, the representation of an entire graph can be obtained. The pooling operation can be defined as *COMBINE* feature as follows:

$$h_u^{(k)} = COMBINE^{(k)}(h_u^{(k)}, a_v^{(k)})$$
(2)

This representation can be then used for various tasks such as classification, prediction, segmentation or reconstruction tasks [6, 8, 14, 17]. The AGGREGATE and COMBINE operators are defined by the specific model.

In this work, we train two separate GNNs: one for H&E graphs and another for TC graphs. Both models are based on a two-layer GNN with an input normalization preprocessing layer. The normalization layer gathers the minimum and maximum values of the input features and, using them, normalizes the features before passing them to the main GNN model.

The number of input features for the H&E model and TC models were 80 and 50, respectively. The hidden layer features of both models were 128, and the output feature after the pooling layer is 1. As the *COMBINE* operator, we adopt a self-attention pooling operation, SAGPool by Lee et al [29]. SAGPool is a hierarchical pooling method that performs local pooling operations of node embeddings in a graph. Each convolutional layer is followed by a SAGPool layer. We used a concatenated jumping knowledge connection [49] to combine the global average of node features obtained from each iteration of graph convolution. This combined representation is passed to a feed-forward network that performs the downstream ordinal classification.



Figure 1. An illustration of the multi-stain workflow using the LateConcatenation architecture. H&E and TC CNNs are used to generate heatmaps from paired H&E and TC WSIs. The latent embedding extracted by H&E and and TC GNNs from the respective graphs are concatenated and a multilayer perceptron predicts a score based on the concatenated embedding.

The output of the feed-forward network is used to learn an ordinal regression model. The goal of an ordinal regression model is to predict labels from an ordinal scale (i.e., labels from a discrete but ordered set) [39]. The decision function used here is Gaussian cumulative link (CL) function which can be defined as:

$$\psi_{CL}(y,\alpha) := \begin{cases} -\log(\sigma(\alpha_1)) & \text{if } y = 1\\ -\log(\sigma(\alpha_y) - \sigma(\alpha_{y-1})) & \text{if } 1 < y < k\\ -\log(1 - \sigma(\alpha_{y-1})) & \text{if } y = k \end{cases}$$
(3)

where  $\sigma$  is the Gaussian cumulative distribution, defined as  $\sigma(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} e^{-x^2/2}$ , k is the number of classes and S is the subset of  $\mathbb{R}^{k-1}$  for which the components are nondecreasing, that is

$$\mathcal{S} := \{ \alpha : \alpha \in \mathbb{R}^{k-1} \text{ and } \alpha_i \leq \alpha_{i+1} \text{ for } 1 \leq i \leq k-1 \}$$

Our GNN performance studies are done using two sets of pathologist labels; those generated by a central pathologist (CP), and those from a group of independent pathologists. The central pathologist scores are from the expert hepatopathologist(s) leading the scoring in the referenced clinical trials. The independent pathologist scores are generated afterwards from hepato-pathologists contracted for this study. One of the challenges with learning from a dataset labeled by a diverse set of raters is the heterogeneity in the systematic per-rater bias present in the labels. This is the case for our independent pathologist scores. There is a bias specific to each pathologist present in these scores.

In order to minimize the effect of pathologist bias on the GNN training, we use a mixed-effect model to learn the pathologist bias and remove it when performing inference [47]. The mixed effect model learns per-pathologist biases in tandem with GNN training. The bias terms are added to the GNN output before mapping to ordinal score values. This minimizes the effect of per-pathologist bias on the GNN training. At inference time, the GNN outputs are directly mapped to ordinal scores.

### 3.4. Fusion

In this section we explore different ways to combine information from the two histological stains. In all the methods listed below, we use the same set of graphs as the ones used to train single stain GNNs. Given paired H&E and TC WSIs for all patients, we aim to learn informative features that summarize the graph representations of WSIs. We also learn a suitable function to merge these features. While previous works [16, 44, 47] have used only one of the two stains to predict fibrosis or NAS component stages, in this work we present a novel multimodal approach that outperforms unimodal baselines when evaluated for concordance with pathologist consensus. We note that the underlying WSIs are not spatially co-registered. The clustering approach to graph generation results in the lack of inter-graph correspondence between H&E and TC graphs. As such, the following methods make no assumptions about the relationship between nodes of the two graphs. Thus, the methods are amenable to other scenarios common in digital pathology where multiple tissue samples are used to make diagnostic or prognostic decisions.

#### 3.4.1 Late Fusion

In this set of methods, information is combined at the very end. Given H&E graph  $\mathcal{H} = (\mathcal{V}_H, \mathcal{E}_H)$  and TC graph  $\mathcal{T} = (\mathcal{V}_T, \mathcal{E}_T)$  and two graph neural networks  $g_H$  and  $g_T$  that operate on  $\mathcal{H}$  and  $\mathcal{T}$  respectively, we obtain vector representation of both graphs as:

$$h = \frac{1}{|\mathcal{V}_H|} \sum_{v \in \mathcal{V}_H} g_H(\mathcal{H}); \ t = \frac{1}{|\mathcal{V}_T|} \sum_{v \in \mathcal{V}_T} g_T(\mathcal{T}) \quad (4)$$

where  $h \in \mathbb{R}^{d_H}$  and  $t \in \mathbb{R}^{d_T}$ ,  $d_H$  and  $d_T$  are the dimensions of the node features in  $\mathcal{H}$  and  $\mathcal{T}$  respectively. Vectors hand t are combined using a combination strategy  $\gamma(.)$  and a feedforward network f maps the multimodal tensor to an ordinal endpoint.

We explore different formulations of  $\gamma(.)$  below:

- Late concatenation:  $\gamma(.)$  takes the form of the concatenation operator and the multimododal tensor  $z \in \mathbb{R}^{d_H+d_T}$ .
- Late addition:  $\gamma(.)$  performs vector addition of latent vectors after projecting them into a common d dimensional latent space as follows:

$$z = W_H \cdot h + W_T \cdot t \tag{5}$$

where  $z \in \mathbb{R}^d$ ,  $W_H$  and  $W_T$  are learnt and  $W_H \in \mathbb{R}^{d \times d_H}$ ,  $W_T \in \mathbb{R}^{d \times d_T}$ .

 Late Hadamard product: γ(.) is the Hadamard product of the projected latent vectors:

$$z = (W_H \cdot h) \circ (W_T \cdot t) \tag{6}$$

where  $z \in \mathbb{R}^d$ ,  $W_H$  and  $W_T$  are learnt and  $W_H \in \mathbb{R}^{d \times d_H}$ ,  $W_T \in \mathbb{R}^{d \times d_T}$ .

• Kronecker product with gating based attention: The fusion strategies discussed above do not explicitly capture interactions across each feature in the latent vector representations of both graphs. Chen et al. [8] use the Kronecker product to fuse information from cell graphs, histology images, and genomic features. They use Kronecker product to ensure every unimodal feature interaction is explicitly accounted for in the multimodal tensor. Their fusion strategy further employs gating-based attention to control the expressiveness of each feature in the unimodal representation. The differentiable gating mechanism learns to suppress the noisy or redundant unimodal feature interactions while retaining the useful interactions. We use this fusion strategy to model interactions between vector representations of  $\mathcal{H}$  and  $\mathcal{T}$ .

$$h' = \alpha_H(h, t) \circ ReLU(W_H \cdot h) \tag{7}$$

$$t' = \alpha_T(h, t) \circ ReLU(W_T \cdot t) \tag{8}$$

where,  $\alpha_m = \sigma(W_{m'} \cdot [h, t])$ 

 $W_m \in \mathbb{R}^{d_m \times d_m}, W_{m'} \in \mathbb{R}^{d_m \times (d_H + d_T)} \quad \forall m \in \{H, T\}$ We learn  $W_m$  and  $W_{m'} \quad \forall m \in \{H, T\}$  for feature gating. Finally, the multimodal tensor is evaluated as:

$$z = \begin{bmatrix} h'\\1 \end{bmatrix} \otimes \begin{bmatrix} t'\\1 \end{bmatrix}$$
(9)

where  $\otimes$  is the Kronecker product and  $z \in \mathbb{R}^{d_H \times d_T}$ . We concatenate 1 to the unimodal vectors to retain unmodified unimodal features in the final multimodal tensor.

### 3.4.2 Mid fusion

Recent work in multimodal deep learning has observed improvements in performance by enabling information to flow across feature encoders much earlier during feature extraction [22,43]. We apply the same intuition and explore methods that enable information to flow between the two graph encoders  $g_H$  and  $g_T$ .

To our knowledge information fusion during graph convolution has not been studied in the literature previously. The work of Fey et al. [12] is similar with some notable differences. In this paper, the authors present a method to pass messages between a molecular graph and its junction tree representation. Intra-messages are passed within a graph along neighboring atoms and inter-messages are passed from atoms of a cluster in the molecular graph to the cluster's junction tree representation and vice versa. Each node in the junction tree represents a meaningful cluster in the original graph e.g, rings or bridged compounds. Inter-messages enable information to be aggregated at the graph level while intra-messages allow for information to flow between the graphs during feature extraction. We note that the junction tree transformation naturally gives rise to node-level correspondence between the molecular graph and its junction tree representation. The authors exploit this characteristic to limit inter-messages to be shared between smaller subgraphs. This is where our problem diverges from that of molecular graphs. In the absence of inter-graph correspondence, we propose two novel strategies to achieve competitive results.

• Global inter message passing (GIMP): Unlike Fey *et al.* [12] we do not have a structured way to pass messages between subgraphs of different modalities. We instead pass a global summary of both graphs to each other at every message passing iteration. Let  $H^{(i)} \in \mathbb{R}^{|V_H| \times d_H}$  and  $T^{(i)} \in \mathbb{R}^{|V_T| \times d_T}$  be the intermediate representations of graphs  $\mathcal{H}$  and  $\mathcal{T}$  after *ith* iteration of graph convolution. We aggregate the node features to obtain the graph summaries  $h^{(i)} \in \mathbb{R}^{d_H}$  and  $t^{(i)} \in \mathbb{R}^{d_T}$  as

$$h^{(i)} = \frac{1}{|V_H|} \sum_{v \in V_H} H^{(i)}; \ t^{(i)} = \frac{1}{|V_T|} \sum_{v \in V_T} T^{(i)}$$
(10)

and pass the graph summaries across the two graphs

$$H_v^{(i)} = H_v^{(i)} + ReLU(W_{T \to H} \cdot t^{(i)}) \; \forall v \in V_H \tag{11}$$

$$T_v^{(i)} = T_v^{(i)} + ReLU(W_{H \to T} \cdot h^{(i)}) \ \forall v \in V_T$$
(12)

where  $W_{T \to H} \in \mathbb{R}^{d_H \times d_T}$  and  $W_{H \to T} \in \mathbb{R}^{d_T \times d_H}$  are learnt.

An iteration of intra-message passing followed by inter-message passing forms one whole iteration of global inter message passing. At the end of graph convolution we concatenate the feature representations obtained from both GNNs similar to late concatenation and use a feedforward network for ordinal classification.

 Global attention inter message passing (GAIMP): The global pooling operation of global inter message passing may severely dilute information for graphs with high cardinality like ours. To overcome this we employ global graph attention proposed by Li *et al.* [31].

$$h^{(i)} = tanh\left(\sum_{v \in V_H} \sigma(f_{H1}(H_v^{(i)})) \circ f_{H2}(H_v^{(i)})\right)$$
(13)

$$t^{(i)} = tanh\left(\sum_{v \in V_T} \sigma(f_{T1}(T_v^{(i)})) \circ f_{T2}(T_v^{(i)})\right)$$
(14)

where  $f_{m1}$  and  $f_{m2} \forall m \in \{H, T\}$  are modality specific attention and projection neural networks respectively. The learned global graph attention helps identify nodes with relevant information which is passed over to all the nodes of the other graph.

## 4. Experimental Setup

The following section describes the data set, evaluation metrics, and other implementation details.

#### 4.1. Data description

Models for segmenting and scoring the NAS components and fibrosis were developed using WSIs from 6 recently completed NASH clinical trials <sup>1</sup>. To evaluate the performance and generalizability of our approach, we used a heldout test dataset which was acquired separate from the clinical trials referred to in the model development set. The experiments in this paper are designed around clinical trial data for two reasons:

 Data heterogeneity Histology data collected from different clinical trials is heterogeneous. This heterogeneity can be attributed to variations in specimen preparation, histology stain, slide scanner and patient populations [23, 41, 45]. For a medical ML system to be useful outside of research, it must perform accurately despite these variations. We find that data heterogeneity is limited in data sets created specifically for research purposes. Campanella *et al.*, 2019 [5] further explore the issue of limited variation in common histopathology datasets like CAMELYON16 [3].

• Data imbalance Real-world medical data typically has class imbalance [30]. This is in contrast to extensively used open-source data sets specifically curated for machine learning research. As before, for a medical ML system to be useful outside of research, it must perform well despite data imbalance.

Slides were scanned at 40X using an Aperio ScanScope® system (Aperio, Vista, CA, USA) and digitized WSIs made available to hepato-pathologists.

### 4.1.1 Patient based data splits

The dataset contains multiple WSIs per patient. For each patient we have one or more H&E and TC slide across different time points. We perform patient-based data splitting to ensure slides from each patient is present in only one of the splits. For the H&E and TC segmentation models, we use 5,175 and 4,737 WSIs in the training set and 1,151 and 1,028 WSIs in the validation set, respectively. The training and the validation sets used for GNN model development consist of 3,734 and 806 pairs of H&E and TC WSIs, respectively. We compare all the GNN models based on their performance on a held-out test set consisting of 620 pairs of H&E and TC WSIs.

#### 4.1.2 Annotations

Board-certified pathologists specializing in hepatobiliary pathology were asked to provide two types of annotation using a digital platform:

• **Pixel-level annotations** To train the CNN models, pixel-level annotations designating regions of tissue within a WSI as exhibiting specific morphologies were used. We collected annotations for 13 H&E classes (bile duct, blood vessels, hepatocellular ballooning, hepatocellular swelling, interface hepatitis, lobular inflammation, lumen, microvesicular steatosis, normal, normal hepatocytes, normal interface, portal inflammation, steatosis) and 5 TC classes (bile duct, blood vessel, fibrosis, lumen, normal). In total, 116,346 annotations of key histological parameters were used for training H&E and TC CNN models.

<sup>&</sup>lt;sup>1</sup>Clinicaltrials.gov identifier: NCT03053050, NCT03053063, NCT01672866, NCT01672879, NCT02466516, NCT03551522; NCT00117676, NCT00116805; NCT01672853

Endpoint	Set	Bin 0	Bin 1	Bin 2	Bin 3	Bin 4
Fibrosis	Dev	239	379	536	1575	2012
	Test	81	402	466	789	95
Hepatocellular ballooning	Dev	990	1301	2878	NA	NA
	Test	319	727	793	NA	NA
Lobular inflammation	Dev	181	1433	1840	1714	NA
	Test	98	1011	664	67	NA
Steatosis	Dev	580	3783	636	184	NA
	Test	81	580	626	553	NA

Table 1. Distribution of NAS component and CRN fibrosis scores for all pathologists in the development and held-out test datasets.

Slide-level annotations Across all the trials in the development set, the biopsy samples were evaluated by a central pathologist who provided slide-level scores based on NASH CRN and NAS scoring system. Additionally, we collected slide-level scores from eight board-certified hepato-pathologists. These pathologists were different from those providing pixel-level annotations. These additional scores allow us to generate a consensus score to compare the ML model against. On the held-out test set, scores were collected from three pathologists. Table 1 presents binwise counts of slide-level annotations used for training and evaluating GNN models.

#### 4.2. Experimental Settings

Pytorch 1.8.1 [38] and Pytorch Geometric 2.0.1 [11] python frameworks were used to train the CNN and GNN models respectively.

Patch-wise CNN Training Using the pathologist annotations as described in section 4.1.3, training image patches on the order of 500,000 samples were generated. These patches were used to train two stain-specific patch-wise CNNs that generate pixel-level predictions of 13 H&E or 5 TC classes. The H&E and TC CNN models are trained separately for a maximum of 30k iterations with validation loss-based early stopping, a batch size of 100 stratified along class and slide ID variables. We use random horizontal flip, random rotation, hue, brightness, contrast, and saturation distortion to augment the patches. We use ADAM optimizer [24] with a learning rate of 1e-4 and learning rate decay factor of 0.5 applied every 10k iterations. Models were composed of 8-12 blocks of compound layers with a topology inspired by residual networks and inception networks with a softmax loss [15, 28].

**GNN Training** We train each GNN for a maximum of 7k iterations with validation loss-based early stopping. We perform an extensive grid search over the number of graph convolution layers, feed-forward layers, hidden dimensions, drop probability, and other model-specific hyperparameters to find the optimal set for each GNN model. The following section summarizes the implementation details of each

GNN model.

**Single Stain GNNs** The two stain-specific GNNs trained are similar in structure with minor differences. Both models use the graph neural network operator proposed by Morris *et al.* [36]. Both models use 2 layers of graph convolution with a SAGPool layer in the middle. The H&E graph nodes have 80 features while the TC graph nodes have 50 features. This difference is due to the distinction in the number of classes present in the H&E and TC heatmaps. The hidden layer dimensions of both GNNs are kept at 128. This is followed by a concatenated jumping connection and feedforward network with 2 layers. We use dropout with a drop probability of 0.5 between the two feed-forward layers.

Late fusion The per-stain graph encoders follow a similar architecture as the single-stain GNNs explained above. In the late fusion models, we remove the last feed-forward layer from GNN architecture and pass the 128-dimensional output vectors to the fusion module. The fusion module for late concatenation, late addition, and late Hadamard product are similar. After performing addition, concatenation, or Hadamard product operations between latent vectors, the multimodal tensor is fed to a three layer feed-forward network using Dropout (drop probability = 0.1) and Batch-Norm. In the case of Kronecker product-based fusion we keep the fusion module architecture close to what is used in the original paper. We use a dropout rate of 0.5 for original architecture and hidden dimension is kept at 64. We show results with both concatenation and bilinear attention computation as described in the paper.

**Mid fusion** The graph encoders in mid-fusion approach follow a similar architecture as the single-stain GNNs. In the GIMP network, we add a projection module to each graph encoder which is composed of a linear layer ( $128 \times 128$ ), ReLU (Rectified Linear Unit) activation and Dropout (drop probability = 0.4). The global attention module of GAIMP uses a linear layer (128,1) followed by sigmoid activation for evaluating attention scores. It uses another linear layer ( $128 \times 128$ ) followed by ReLU activation and Dropout (drop probability = 0.2) for projection.

#### 4.3. Evaluation

We use linearly weighted Cohen's kappa to evaluate model performance on the held-out test set. Cohen's kappa is a form of accuracy metric that corrects for agreement by chance amongst two raters [34]. Disagreements are penalized proportional to the absolute difference between the two ratings [42]. We apply bootstrap analysis (n = 400) to the kappa calculation to study the effects of sampling variations for model and inter-pathologist performance.

## 5. Results and Discussion

We compare the GNN models using scores from all pathologists. Models learn from more than one scores per

Model/label	Ballooning	Lobular inflammation	Steatosis	Fibrosis
H&E [unimodal]	0.49 [0.43,0.53]*	0.43 [0.37,0.47]*	0.69 [0.65,0.73]*	0.48 [0.43,0.52]
Trichrome [unimodal]	0.27 [0.22,0.33]	0.23 [0.18,0.28]	0.16 [0.12,0.2]	0.51 [0.47,0.56]*
GAIMP [fusion]	0.41 [0.36,0.45]	0.37 [0.32,0.41]	0.7 [0.66,0.74]	0.58 [0.54,0.62]
GIMP [fusion]	0.46 [0.41,0.5]	0.41 [0.35,0.46]	0.67 [0.62,0.71]	0.56 [0.52,0.6]
LateConcatenation [fusion]	0.52 [0.47,0.57]	0.42 [0.38,0.48]	0.67 [0.63,0.72]	<b>0.61</b> [0.57,0.65]
LateAddition [fusion]	<b>0.54</b> [0.49,0.58]	0.45 [0.4,0.5]	0.7 [0.66,0.74]	0.6 [0.56,0.64]
LateHadamardProduct [fusion]	0.37 [0.32,0.41]	<b>0.51</b> [0.45,0.56]	0.63 [0.59,0.67]	0.48 [0.44,0.53]
pathomicFusion [fusion]	0.53 [0.48,0.58]	0.41 [0.37,0.46]	0.7 [0.66,0.74]	0.59 [0.55,0.63]
pathomicFusionBilinear [fusion]	0.45 [0.41,0.5]	0.4 [0.35,0.44]	<b>0.71</b> [0.67,0.75]	0.6 [0.55,0.63]
Pathologist	0.49 [0.42,0.6]	0.36 [0.27, 0.45]	0.6 [0.49,0.69]	0.54 [0.34,0.66]

Table 2. Linearly weighted Cohen's kappa is used to compare performance of fusion multi-stain models, standard single-stain models, and pathologists measured against consensus pathologist scores

WSI by means of a mixed-effect model which accounts for heterogeneity in pathologist bias [40]. The learned biases are discarded during inference as described in Section 3.3.

In Table 2, H&E and Trichrome rows at the top represent the single-stain GNNs trained on unimodal graphs. The asterisks<sup>\*</sup> indicate the primary single-stain model for a given endpoint. H&E is the primary stain for scoring ballooning, lobular inflammation, and steatosis by pathologists, and TC is the primary stain for scoring fibrosis stage. GIMP and GAIMP in Table 2 refer to global inter message passing and global attention inter message passing, respectively. For the models, we report the mean linearly weighted Cohen's kappa measured between the model and median consensus of all pathologists [2] from bootstrap analysis in addition to the 95% confidence interval around the mean. Similarly, in the last row of Table 2, we report the mean linearly weighted Cohen's kappa for independent pathologists as a performance baseline. This kappa measures agreement between independent pathologist scores and consensus (median) pathologist scores. The performance of single-stain ML models (rows 1 and 2 in Table 2) match independent pathologists, which is in agreement with prior published work [44].

Based on results on the test set shown in Table 2, the LateConcatenation, LateAddition, and pathomicFusion approaches either outperform or perform on par with the single-stain models. Most notably, these three fusion approaches statistically outperform unimodal (single-stain) H&E-based prediction of the ballooning score and statistically outperform unimodal TC-based prediction and pathologist scoring of the CRN fibrosis score. These improvements in performance suggest the presence of complementary and independently predictive information in both stains.

Further, in Table 2 we report a maximum relative improvement of 20% between LateConcatenation and TC, which is the primary unimodal model in the prediction of fibrosis. Similarly, we observe a 19% relative improvement in lobular inflammation for LateHadamardProduct, 10% in

ballooning for LateAddition, and 3% in steatosis for pathomicFusionBilinear over prediction from H&E, which is the primary unimodal model for prediction of NAS components. Overall, LateAddition fusion model provides the most consistent performance with 18% relative improvement in fibrosis, 10% in ballooning, 5% in lobular inflammation and 1% in steatosis.

The improved performance observed in fusion models over unimodal single-stain models can be primarily attributed to the richness of information present in the histologic stains in relation to their non-primary endpoints, e.g., H&E for fibrosis. In samples with sub-optimal quality or inadequate tissue in the primary WSI, the fusion model can learn to use information from the non-primary stain, e.g., leverage H&E when TC staining is poor, to provide a more accurate score. This is an advantage that fusion models can gain over unimodal models from having access to multimodal data. It must be noted that in NASH pathology it is common to leverage the non-primary stain when the primary stain has quality complications. It has also been reported in recent work [26] that models can locate and segment non-primary features such as collagen in H&E samples which are aligned with observations in this work.

## 6. Conclusions

This paper presents a novel graph fusion-based approach for evaluating NAS components and CRN fibrosis stage given paired H&E and TC histology slides. Both H&E and TC slides are collected from patients under examination for NASH. Pathologists and previous ML systems typically use one (primary) of the two slides to evaluate either the NAS component scores or the CRN fibrosis score. In this paper, we show that using additional information in the non-primary whole-slide image improves the performance of models predicting NASH scores substantially. We show that our fusion strategy statistically outperforms unimodal ML models and pathologists in at least two NASH score components on held-out clinical data.

## References

- [1] Mohammed Adnan, Shivam Kalra, and Hamid R Tizhoosh. Representation learning of histopathology images using graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 988–989, 2020. 1, 2
- [2] David Ahmedt-Aristizabal, Mohammad Ali Armin, Simon Denman, Clinton Fookes, and Lars Petersson. A survey on graph-based deep learning for computational histopathology. *Computerized Medical Imaging and Graphics*, 95:102027, 2022. 1
- [3] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes van Diest, Bram van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen A. W. M. van der Laak, Meyke Hermsen, Quirine F Manson, Maschenka Balkenhol, Oscar Geessink, Nikolaos Stathonikos, Marcory CRF van Dijk, Peter Bult, Francisco Beca, Andrew H Beck, Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, Aoxiao Zhong, Qi Dou, Quanzheng Li, Hao Chen, Huang-Jing Lin, Pheng-Ann Heng, Christian Haß, Elia Bruni, Quincy Wong, Ugur Halici, Mustafa Ümit Öner, Rengul Cetin-Atalay, Matt Berseth, Vitali Khvatkov, Alexei Vylegzhanin, Oren Kraus, Muhammad Shaban, Nasir Rajpoot, Ruqayya Awan, Korsuk Sirinukunwattana, Talha Qaiser, Yee-Wah Tsang, David Tellez, Jonas Annuscheit, Peter Hufnagl, Mira Valkonen, Kimmo Kartasalo, Leena Latonen, Pekka Ruusuvuori, Kaisa Liimatainen, Shadi Albarqouni, Bharti Mungal, Ami George, Stefanie Demirci, Nassir Navab, Seiryo Watanabe, Shigeto Seno, Yoichi Takenaka, Hideo Matsuda, Hady Ahmady Phoulady, Vassili Kovalev, Alexander Kalinovsky, Vitali Liauchuk, Gloria Bueno, M. Milagro Fernandez-Carrobles, Ismael Serrano, Oscar Deniz, Daniel Racoceanu, and Rui Venâncio and. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. JAMA, 318(22):2199, Dec. 2017. 6
- [4] Gregory Thomas Brown and David E Kleiner. Histopathology of nonalcoholic fatty liver disease and nonalcoholic steatohepatitis. *Metabolism*, 65(8):1080–1086, 2016. 1
- [5] Gabriele Campanella, Matthew G. Hanna, Luke Geneslaw, Allen Miraflor, Vitor Werneck Krauss Silva, Klaus J. Busam, Edi Brogi, Victor E. Reuter, David S. Klimstra, and Thomas J. Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine*, 25(8):1301–1309, July 2019. 6
- [6] Chun-Hung Chao, Zhuotun Zhu, Dazhou Guo, Ke Yan, Tsung-Ying Ho, Jinzheng Cai, Adam P Harrison, Xianghua Ye, Jing Xiao, Alan Yuille, et al. Lymph node gross tumor volume detection in oncology imaging via relationship learning using graph neural network. In *International Conference* on Medical Image Computing and Computer-Assisted Intervention, pages 772–782. Springer, 2020. 1, 3
- [7] Linyan Chen, Hao Zeng, Yu Xiang, Yeqian Huang, Yuling Luo, and Xuelei Ma. Histopathological images and multiomics integration predict molecular characteristics and survival in lung adenocarcinoma. *Frontiers in Cell and Devel*opmental Biology, 9, Oct. 2021. 2

- [8] Richard J Chen, Ming Y Lu, Jingwen Wang, Drew FK Williamson, Scott J Rodig, Neal I Lindeman, and Faisal Mahmood. Pathomic fusion: an integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. *IEEE Transactions on Medical Imaging*, 2020. 2, 3, 5
- [9] Richard J. Chen, Ming Y. Lu, Wei-Hung Weng, Tiffany Y. Chen, Drew F.K. Williamson, Trevor Manz, Maha Shady, and Faisal Mahmood. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4025, October 2021. 2
- [10] Thomas G. Dietterich, Richard H. Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial Intelligence*, 89(1):31–71, 1997. 2
- [11] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric, 2019. 7
- [12] Matthias Fey, Jan-Gin Yuen, and Frank Weichert. Hierarchical inter-message passing for learning on molecular graphs. arXiv preprint arXiv:2006.12179, 2020. 5
- [13] Shrey Gadiya, Deepak Anand, and Amit Sethi. Histographs: Graphs in histopathology, 2019. 2
- [14] Karthik Gopinath, Christian Desrosiers, and Herve Lombaert. Graph convolutions on spectral embeddings for cortical surface parcellation. *Medical image analysis*, 54:297– 305, 2019. 3
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning. *Image Recognition*, 2015. 7
- [16] Fabian Heinemann, Gerald Birk, and Birgit Stierstorfer. Deep learning enables pathologist-like scoring of NASH models. *Scientific Reports*, 9(1), Dec. 2019. 1, 2, 4
- [17] Yoonmi Hong, Geng Chen, Pew-Thian Yap, and Dinggang Shen. Multifold acceleration of diffusion mri via deep learning reconstruction from slice-undersampled data. In *International Conference on Information Processing in Medical Imaging*, pages 530–541. Springer, 2019. 3
- [18] Osamu Iizuka, Fahdi Kanavati, Kei Kato, Michael Rambeau, Koji Arihiro, and Masayuki Tsuneki. Deep learning models for histopathological classification of gastric and colonic epithelial tumours. *Scientific Reports*, 10(1), Jan. 2020. 2
- [19] Maximilian Ilse, Jakub M. Tomczak, and Max Welling. Chapter 22 - deep multiple instance learning for digital histopathology. In S. Kevin Zhou, Daniel Rueckert, and Gabor Fichtinger, editors, *Handbook of Medical Image Computing and Computer Assisted Intervention*, The Elsevier and MICCAI Society Book Series, pages 521–546. Academic Press, 2020. 2
- [20] National Cancer Institute. How cancer is diagnosed. 2
- [21] Andrew Janowczyk and Anant Madabhushi. Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases. *Journal of Pathology Informatics*, 7(1):29, 2016. 2
- [22] Hamid Reza Vaezi Joze, Amirreza Shaban, Michael L. Iuzzolino, and Kazuhito Koishida. MMTM: multimodal transfer module for CNN fusion. *CoRR*, abs/1911.08670, 2019. 5

- [23] Christopher J. Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17(1), Oct. 2019. 6
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. 7
- [25] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net, 2017. 3
- [26] Pawel Kleczek, Joanna Jaworek-Korjakowska, and Marek Gorgon. A novel method for tissue segmentation in high-resolution h&e-stained histopathological whole-slide images. *Computerized Medical Imaging and Graphics*, 79:101686, 2020. 8
- [27] Daisuke Komura and Shumpei Ishikawa. Machine learning methods for histopathological image analysis. *Computational and structural biotechnology journal*, 16:34–42, 2018.
   2
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25, 2012. 2, 7
- [29] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. Self-attention graph pooling. In *International conference on machine learning*, pages 3734–3743. PMLR, 2019. 3
- [30] Der-Chiang Li, Chiao-Wen Liu, and Susan C. Hu. A learning method for the class imbalance problem with medical data sets. *Computers in Biology and Medicine*, 40(5):509–518, 2010. 6
- [31] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S. Zemel. Gated graph sequence neural networks. In Yoshua Bengio and Yann LeCun, editors, 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings, 2016. 6
- [32] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 3
- [33] Anant Madabhushi and George Lee. Image analysis and machine learning in digital pathology: Challenges and opportunities. *Medical image analysis*, 33:170–175, 2016. 2
- [34] Mary L McHugh. Interrater reliability: the kappa statistic. Biochem. Med. (Zagreb), 22(3):276–282, 2012. 7
- [35] Pooya Mobadersany, Safoora Yousefi, Mohamed Amgad, David A. Gutman, Jill S. Barnholtz-Sloan, José E. Velázquez Vega, Daniel J. Brat, and Lee A. D. Cooper. Predicting cancer outcomes from histology and genomics using convolutional networks. *Proceedings of the National Academy of Sciences*, 115(13), Mar. 2018. 2
- [36] Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin

Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. *CoRR*, abs/1810.02244, 2018. 7

- [37] Raouf E. Nakhleh, Vania Nosé, Carol Colasacco, Lisa A. Fatheree, Tamera J. Lillemoe, Douglas C. McCrory, Frederick A. Meier, Christopher N. Otis, Scott R. Owens, Stephen S. Raab, Roderick R. Turner, Christina B. Ventura, and Andrew A. Renshaw. Interpretive diagnostic error reduction in surgical pathology and cytology: Guideline from the college of american pathologists pathology and laboratory quality center and the association of directors of anatomic and surgical pathology. Archives of Pathology & Laboratory Medicine, 140(1):29–40, May 2015. 2
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. 7
- [39] Fabian Pedregosa, Francis Bach, and Alexandre Gramfort. On the consistency of ordinal regression methods. *Journal of Machine Learning Research*, 18:1–35, 2017. 4
- [40] Omid Pournik, Seyed Moayed Alavian, Leila Ghalichi, Bahram Seifizarei, Leila Mehrnoush, Azam Aslani, Soghra Anjarani, and Saeid Eslami. Inter-observer and intraobserver agreement in pathological evaluation of nonalcoholic fatty liver disease suspected liver biopsies. *Hepatitis monthly*, 14(1), 2014. 8
- [41] Prarthana Shrestha, R Kneepkens, J Vrijnsen, D Vossen, E Abels, and B Hulsken. A quantitative approach to evaluate image quality of whole slide imaging scanners. J. Pathol. Inform., 7(1):56, Dec. 2016. 6
- [42] Julius Sim and Chris C Wright. The kappa statistic in reliability studies: use, interpretation, and sample size requirements. *Physical therapy*, 85(3):257–268, 2005. 7
- [43] Hao Tan and Mohit Bansal. LXMERT: learning crossmodality encoder representations from transformers. *CoRR*, abs/1908.07490, 2019. 5
- [44] Amaro Taylor-Weiner, Harsha Pokkalla, Ling Han, Catherine Jia, Ryan Huss, Chuhan Chung, Hunter Elliott, Benjamin Glass, Kishalve Pethia, Oscar Carrasco-Zevallos, et al. A machine learning approach enables quantitative measurement of liver histology and disease monitoring in nash. *Hepatology*, 74(1):133–147, 2021. 1, 2, 3, 4, 8
- [45] David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen van der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical Image Analysis*, 58:101544, 2019. 6
- [46] Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H. Beck. Deep learning for identifying metastatic breast cancer, 2016. 2
- [47] Jason Wang, Maryam Pouryahya, Kenneth Leidal, Harsha Pokkalla, Dinkar Juyal, Zahil Shanis, Aryan Pedawi, Quang Le, Victoria Mountain, Sara Hoffman, et al. Liver biopsy

graph neural networks for automated histologic scoring using the nash crn system. In *JOURNAL OF HEPATOLOGY*, volume 75, pages S602–S603. ELSEVIER RADARWEG 29, 1043 NX AMSTERDAM, NETHERLANDS, 2021. 1, 2, 3, 4

- [48] Christopher D Williams, Joel Stengel, Michael I Asike, Dawn M Torres, Janet Shaw, Maricela Contreras, Cristy L Landt, and Stephen A Harrison. Prevalence of nonalcoholic fatty liver disease and nonalcoholic steatohepatitis among a largely middle-aged population utilizing ultrasound and liver biopsy: a prospective study. *Gastroenterology*, 140(1):124– 131, 2011. 1
- [49] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *International Conference on Machine Learning*, pages 5453–5462. PMLR, 2018. 3
- [50] Zobair M Younossi, Deirdre Blissett, Robert Blissett, Linda Henry, Maria Stepanova, Youssef Younossi, Andrei Racila, Sharon Hunt, and Rachel Beckerman. The economic and clinical burden of nonalcoholic fatty liver disease in the united states and europe. *Hepatology*, 64(5):1577–1586, 2016. 1
- [51] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. Birch: an efficient data clustering method for very large databases. *ACM sigmod record*, 25(2):103–114, 1996. 3