

Multimodal Optimal Transport-based Co-Attention Transformer with Global Structure Consistency for Survival Prediction

Yingxue Xu Hao Chen*

The Hong Kong University of Science and Technology

{yxueb, jhc}@cse.ust.hk

Abstract

Survival prediction is a complicated ordinal regression task that aims to predict the ranking risk of death, which generally benefits from the integration of histology and genomic data. Despite the progress in joint learning from pathology and genomics, existing methods still suffer from challenging issues: 1) Due to the large size of pathological images, it is difficult to effectively represent the gigapixel whole slide images (WSIs). 2) Interactions within tumor microenvironment (TME) in histology are essential for survival analysis. Although current approaches attempt to model these interactions via co-attention between histology and genomic data, they focus on only dense local similarity across modalities, which fails to capture global consistency between potential structures, i.e. TME-related interactions of histology and co-expression of genomic data. To address these challenges, we propose a Multimodal Optimal Transport-based Co-Attention Transformer framework with global structure consistency, in which optimal transport (OT) is applied to match patches of a WSI and genes embeddings for selecting informative patches to represent the gigapixel WSI. More importantly, OT-based co-attention provides a global awareness to effectively capture structural interactions within TME for survival prediction. To overcome high computational complexity of OT, we propose a robust and efficient implementation over micro-batch of WSI patches by approximating the original OT with unbalanced mini-batch OT. Extensive experiments show the superiority of our method on five benchmark datasets compared to the state-of-the-art methods. The code is released¹.

1. Introduction

Survival prediction is a complex ordinal regression task that aims to estimate the relative risk of death in cancer prognosis, which generally integrates the manual as-

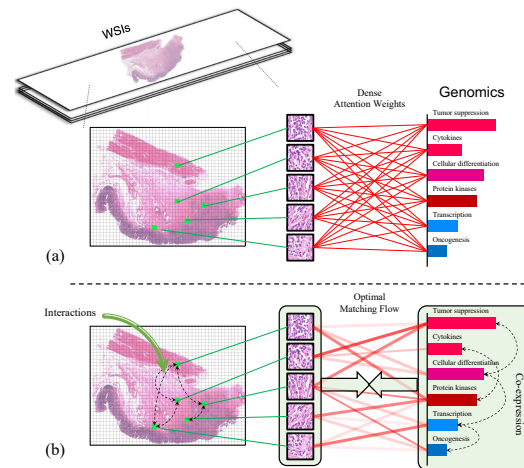


Figure 1. Comparison between (a) Co-Attention and (b) OT-based Co-Attention, where (a) learns weights of instances by dense attention, and (b) identifies informative instances by optimal transport matching flow from a global perspective, which enforces considering potential structures of each modality, i.e. interactions within WSIs and co-expression within genomics.

essment of qualitative morphological information from pathology data and quantitative molecular profiles from genomic data in clinical settings [6, 55, 4]. Despite recent advances in multimodal learning for histology and genomics [5, 55, 27, 34], there still exists several open issues. Among others, one daunting challenge is to capture key information from heterogeneous modalities for effective fusion. Particularly, due to the large size of about $500,000 \times 500,000$ pixels, it is challenging to effectively represent gigapixel whole slide images without losing key information. Furthermore, visual concepts of tumor microenvironment (TME) within pathological images are verified to have significant associations with survival analysis in various cancer types [42, 1, 25], e.g. cellular components including fibroblast cells and various immune cells that can alter cancer cell behaviors [33]. However, the TME-related patches occupy only a tiny proportion of the entire WSI, leading to a fine-grained visual recognition problem [54] that is indiscernible

*Corresponding author

¹<https://github.com/Innse/MOTCat>

by conventional multimodal learning.

Recently attention-based multiple instance learning (MIL) [19, 26] has provided a typical solution to identify informative instances, in which a WSI is formulated as a bag of patches (instances) and an attention score is assigned for each instance as a weight for selection. In multimodal learning, genomic data has been applied to guide the selection of TME-related instances by co-attention mechanism across modalities [5], as genes expression might correspond to some morphological characteristics revealed in pathological TME [29, 38]. They [5] densely calculate similarity scores for each pair of pathology and genomic instances as weights of selection to capture fine-grained visual information in WSIs, as shown in Fig. 1 (a). However, this type of approaches with a local view may not thoroughly learn information about TME, since they ignore global potential structure [55] within modality, e.g. interactions within TME for histology and co-expression for genomics [31, 51].

Fruitful works [31, 33, 5] demonstrate that the interactions within TME are important indicators affecting survival outcomes, e.g. co-occurrence of tumor cells with tumor-infiltrating lymphocytes (TILs) of the immune system is a positive prognostic indicator. However, these collaborative components in TME might be spatially dispersed throughout the entire WSI, which indicates long-range structural associations in WSI, as shown in Fig. 1 (b). On the other hand, genes co-expression [49, 55] also suggests a potential structure. There might be intrinsic consistency between these two potential structures [55, 31], as some studies have verified that genomic mutations can alter normal functions and biological processes of TILs within TME [51, 47]. By leveraging the global consistency between their potential structures to match histology and genomic data, it is more likely to identify TME-related patches from WSIs. However, existing co-attention-based multimodal learning focuses on only dense local similarity, neglecting the global coherence of potential structures.

To address these challenges, we propose a **Multimodal Optimal Transport-based Co-Attention Transformer (MOTCat)** framework with global structure consistency, in which optimal transport-based co-attention is applied to match instances between histology and genomics from a global perspective, as shown in Fig. 1 (b). Optimal transport (OT) [11, 3, 41], as a structural matching approach, is able to generate an optimal matching solution with the overall minimum matching cost, based on local pairwise cost between instances of histology and genomics. As a result, instances of a WSI that have high global structure consistency with genes co-expressions can be identified to represent the WSI. These instances might be more strongly associated with TME that contributes to survival prediction. In this way, the aforementioned two challenges can be addressed by OT-based Co-Attention simultaneously.

Compared with the conventional co-attention mechanism, the advantages of the proposed OT-based co-attention are three-fold. First, optimal transport provides the instances matching process with global awareness. Marginal constraints of total mass equality enforce a trade-off of instances within modality for transporting during optimization, instead of only considering the local similarity of pairwise instances as conventional co-attention does. It enables the modeling of structural interactions of WSIs and co-expressions of genomics, which is beneficial to survival prediction. Second, the learned patch-level attention score is not a rigorous metric for selecting informative instances under weak supervision [48], i.e. slide-level survival time, especially with only a small number of WSI samples. As an alternative, the optimal matching flow is the rigorous closed-form solution to OT problem without accessing any label, achieved by solving the Linear Programming problem. Last, the optimal matching flow provides a transformation between two modalities under preserving the potential structure, which reduces cross-modal heterogeneity gap.

Nevertheless, due to a massive number of patches from gigapixel WSIs, OT-based co-attention might suffer from high computational complexity, preventing from applying OT to pathology data in practice. To address this, we propose a robust and efficient implementation of OT-based co-attention for matching histology and genomics. Specifically, we split all instances of a WSI into a subset termed **Micro-Batch** and get the averaged result as a proxy to approximate the solution to the original OT problem over all instances by unbalanced optimal transport [13], which can provide a more robust approximation to the subset sampling with the theoretical and experimental guarantee.

It is worth noting that the proposed method is a generalized multimodal learning framework that captures potential structure across modalities with global structure consistency. The contributions of this work are summarized as follows: (1) we propose a novel multimodal OT-based Co-Attention Transformer with global structure consistency, where OT-based co-attention is used to identify informative instances by global optimal matching, which allows modeling interactions of histology and co-expression of genomics. (2) We propose a robust and efficient implementation of OT-based co-attention over **Micro-Batch**. (3) Extensive experiments on five benchmark datasets show significant improvement over state-of-the-art methods.

2. Related Work

2.1. Multimodal Learning in Healthcare

In clinical settings, the patient state is often characterized by a spectrum of various modalities [30], such as pathology [54, 32], radiology [39, 50], genomics [28, 47, 31], etc. Multimodal learning has shown superiority by leveraging

the complementary information from multimodal data [30], in which the key issue is to overcome data heterogeneity for better feature fusion. Most existing methods can be roughly classified into three categories: early fusion [18, 43], late fusion [20, 4], and intermediate fusion [5, 24]. Early fusion is to aggregate all modalities at the input level. The most straightforward solution is fusion operators such as concatenation [18], Kronecker Product [35], etc. However, it neglects intra-modality dynamics [30]. On the contrary, late fusion [15] integrates the predictions from individually separated models at the decision level for the final decision, which cannot fully explore cross-modal interactions. Recently, intermediate fusion has attracted much interest by capturing cross-modal interconnections at different levels, where the typical one is attention-based multimodal fusion [5]. For example, Zhu et al. [52] proposed a triplet attention to integrate MRI and diffusion tensor imaging (DTI) for epilepsy diagnosis. IMGFN [55] designed a graph attention network for survival prediction. MCAT [5] proposed a co-attention to identify informative instances of gigapixel WSI with genomic features as queries. HFBSurv [27] progressively integrated multimodal information from the low level to the high level for survival analysis. Note that our method MOTCat belongs to one of the intermediate fusion approaches, which explores the global consistency of potential structure by optimal transport.

2.2. Multiple Instance Learning in Pathology

Since it is difficult to represent WSIs of large size, MIL algorithms have achieved remarkable progress in pathology, in which WSIs are often formulated as a bag of pathology patches. MIL in pathology generally can be divided into two categories: instance-based algorithms [2, 44, 7] and embedding-based algorithms [19, 16, 32]. The former aims to select a small number of instances within each WSI for training an aggregation model. The latter one is to map each instance into a low-dimension fixed-length embedding offline, and then to learn a bag-level representation based on instance-level embeddings. There are several strategies for aggregating instance-level embeddings. One is clustering-based methods [37, 40], where all the instance-level embeddings are clustered to several centroids that are integrated for the final prediction. The most common strategy is attention-based MIL (AB-MIL) [19, 26] that assigns a weight for each instance, in which various approaches are developed differing in the ways to generate the weighting values. For example, one of the first AB-MIL [19] used a side-branch network for generating attention weights, followed by DS-MIL [26] adopted cosine distance to measure the similarity of instances as weights. Recently, transformer-based MIL like TransMIL [36] has leveraged the self-attention mechanism to explore long-range interactions in WSI. Furthermore, DTFD-MIL [48]

proposed the double-tier MIL framework for more rigorous instance weights. Unlike AB-MIL, we use optimal match flow generated by OT between pathological and genomic instances to identify instances of WSI that have global potential structure consistency with that of genomics.

2.3. Survival Prediction

Survival outcome prediction, also known as time-to-event analysis [23], aims to predict the probability of experiencing an event of interest, e.g. death event in the clinical setting, before a specific time under both right-censored and uncensored data. Right-censored data refers to cases where the outcome event cannot be observed within the study period [23]. Before the deep learning era, statistical models dominated survival analysis. The typical one is Cox Proportional Hazards (CoxPH) model [9] that characterizes the hazard function, the risk of death, by an exponential linear model. DeepSurv [12] is the first work incorporating deep learning into survival analysis, which models the risk function by several fully connected layers and feeds it to the CoxPH model for the final hazard function. Furthermore, DeepConvSurv [53] first attempted to use pathological images for a deep survival model. Recent works [27, 6, 34, 4, 5] tend to directly estimate the hazard function by deep learning models without any statistical assumption. Histology and genomics data are often integrated as the gold standard for predicting survival outcomes with promising performance [4, 6, 55]. Similarly, our method incorporates these two modalities for better survival prediction by learning their global structural coherence.

3. Method

3.1. Overview and Problem Formulation

The overall framework of the proposed method is shown in Fig. 2. In this work, we aim to estimate the hazard functions $h_n(t)$ considered as the risk probability of experiencing death event in a small interval after time point t for n -th patient, based on pathology data \mathbf{X}_n^p and genomic data \mathbf{X}_n^g .

For pathology, we formulate each WSI as a "bag" structure in a conventional multiple instance learning (MIL) setting, where a bag consists of a collection of instances that are patched from the tissue-containing image regions of a WSI, and instance-level features are extracted from patch images. Genomic data is also organized as a "bag" based on the biological functional impact of genes. For each instance of the genomic bag, we adopt an encoder to extract genes embeddings. The bags formulations for histology and genomics are illustrated in Sec. 3.1.1 and 3.1.2, respectively.

Then, Optimal Transport-based Co-Attention is applied to match instances of a WSI and the corresponding genomic instances to model pathological interactions and genomic co-expressions, so that the instances with high structure

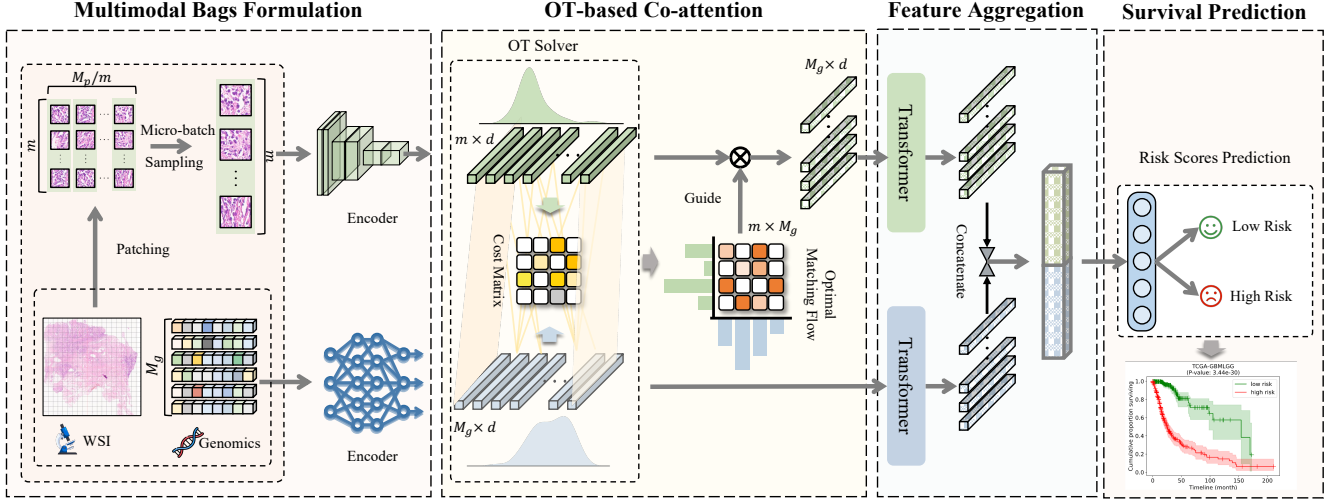


Figure 2. Overview of Multimodal Optimal Transport-based Co-Attention Transformer (MOTCat) architecture, in which both modalities are formulated as bags representations and a micro-batch of the WSI bag is sampled as pathological input. Then OT-based Co-attention is to calculate optimal matching flow for identifying informative instances with global structure consistency. After aggregating selected instances, features from both modalities are concatenated for survival prediction.

consistency can be identified to represent the whole modality. This part will be demonstrated in Sec. 3.2.

After that, a Transformer encoder will integrate all instances after selection for each modality into a bag-level representation, in which each instance is treated as a token. The bag-level representation will be concatenated into the multimodal features for predicting a hazard function to get the ordinal survival risk. The details will be described in Sec. 3.1.3. Finally, optimization process over Micro-Batch and the corresponding computational complexity analysis will be presented in Sec. 3.3.

3.1.1 WSI Bags Formulation

We formulate the learning process of pathology images (WSIs) as a weakly-supervised MIL task, in which each WSI \mathbf{X}_n^p is represented as a bag with access to only bag-level (slide-level) labels. Given a bag $\mathbf{X}_n^p = \{\mathbf{x}_{n,i}^p\}_{i=1}^{M_p}$ of M_p permutation-invariant instances, the goal is to encode instance-level features of a WSI into a compact bag-level feature embedding and then assign a bag-level label to it. To this end, we apply a CNN encoder $f_p(\cdot)$ for each patch of the WSI to obtain a bag \mathbf{B}_n^p of instance-level features:

$$\mathbf{B}_n^p = \{f_p(\mathbf{x}_{n,i}^p) : \mathbf{x}_{n,i}^p \in \mathbf{X}_n^p\} = \{\mathbf{b}_{n,i}^p\}_{i=1}^{M_p}, \quad (1)$$

where $\mathbf{B}_n^p \in \mathbb{R}^{M_p \times d}$ contains d -dimensional instances and $\mathbf{b}_{n,i}^p = f_p(\mathbf{x}_{n,i}^p)$. After identifying informative instances of \mathbf{B}_n^p by OT-based Co-Attention, we construct a new bag $\hat{\mathbf{B}}_n^p \in \mathbb{R}^{\hat{M}_p \times d}$ of features with \hat{M}_p instances, which will be illustrated in detail in Section 3.2. Then, a transformer encoder is used as a global aggregation model $\mathcal{T}_p(\cdot)$ to obtain

the bag-level representation $H_n^p = \mathcal{T}_p(\hat{\mathbf{B}}_n^p)$ for the subsequent features fusion.

Algorithm 1 The MOTCat Algorithm

Input: A WSI bag as $\mathbf{X}_n^p = \{\mathbf{x}_{n,i}^p\}_{i=1}^{M_p}$; A genomic bag as $\mathbf{X}_n^g = \{\mathbf{x}_{n,j}^g\}_{j=1}^{M_g}$. m refers to the size for a micro-batch of a WSI.

Output: The optimal parameters of model;

- 1: // Sample micro-batch from the WSI bag $X_n^{p,m} = \{\mathbf{x}_{n,i}^p\}_{i=1}^m$ each time.
 - 2: **for** $X_n^{p,m}$ in \mathbf{X}_n^p **do**
 - 3: Fix model parameters f_p, f_g, \mathcal{T}_p and \mathcal{T}_g .
 - 4: // To extract a bag of instance-level features
 - 5: // Eq. 1 over micro-batch of histology
 - 6: $B_n^{p,m} = \{f_p(\mathbf{x}_{n,i}^p) : \mathbf{x}_{n,i}^p \in X_n^{p,m}\}$
 - 7: // Eq. 2 of genomics
 - 8: $\mathbf{B}_n^g = \{f_g^j(\mathbf{x}_{n,j}^g) : \mathbf{x}_{n,j}^g \in \mathbf{X}_n^g\}$
 - 9: Optimize \mathbf{P}_n^m via Sinkhorn-Knopp matrix scaling algorithm [8, 14]:
 - 10: $\mathbf{P}_n^m \leftarrow \arg \min_{\mathbf{P}_n^m} \mathcal{W}^m(B_n^{p,m}, \mathbf{B}_n^g)$ // Eq. 5
 - 11: Fix \mathbf{P}_n^m and select instances from $B_n^{p,m}$ by:
 - 12: $\hat{B}_n^{p,m} = \mathbf{P}_n^{m \top} B_n^{p,m}$
 - 13: // To get bag-level representation
 - 14: $H_n^{p,m} = \mathcal{T}_p(\hat{B}_n^{p,m})$ // for histology
 - 15: $H_n^g = \mathcal{T}_g(\hat{\mathbf{B}}_n^g)$ // for genomics
 - 16: $H_n^m = H_n^{p,m} \bowtie H_n^g$ // Concatenation
 - 17: Calculate the overall loss function \mathcal{L} of Eq. 6.
 - 18: Update parameters of $f_g, \mathcal{T}_p, \mathcal{T}_g$ and the last fully connected layers of f_p by Adam.
 - 19: **end for**
-

3.1.2 Genomic Bags Formulation

Genomic data consists of 1×1 attributes, such as mutation status, transcript abundance (RNA-Seq), copy number variation (CNV) and other molecular characterizations. Following the organization of genomic data in the previous work [5, 28] that categories them according to the biological functional impact, the expressive genomic bag can be organized as $\mathbf{X}_n^g = \{\mathbf{x}_{n,j}^g\}_{j=1}^{M_g} \in \mathbb{R}^{M_g \times d_j}$ with M_g instances of unique functional categories. For j -th category, the corresponding genomic embedding consists of d_j -dimensional attributes encoded by a separated network $f_g^j(\cdot)$. Similarly, a bag $\mathbf{B}_n^g \in \mathbb{R}^{M_g \times d}$ of genomic data with M_g instances can be constructed as follows:

$$\mathbf{B}_n^g = \{f_g^j(\mathbf{x}_{n,j}^g) : \mathbf{x}_{n,j}^g \in \mathbf{X}_n^g\} = \{\mathbf{b}_{n,j}^g\}_{j=1}^{M_g}, \quad (2)$$

which is aggregated to the bag-level embedding by a transformer encoder $H_n^g = \mathcal{T}_g(\mathbf{B}_n^g)$ as well.

3.1.3 Multimodal Survival Prediction Formulation

Survival outcome prediction aims to predict the risk probability of an outcome event occurring before a specific time, in which the outcome event is not always observed, leading to a right-censored event. In this task, let $c \in \{0, 1\}$ be censor status indicating whether the outcome event is right-censored ($c = 1$) or not ($c = 0$), and $t \in \mathbb{R}^+$ refers to the overall survival time (in months). Given a training set of N pathology-genomics pairs $\mathcal{D} = \{(\mathbf{X}_n^p, \mathbf{X}_n^g), c_n, t^n\}_{n=1}^N$, we can acquire the bag-level features H_n^p of WSIs and H_n^g of genomic data for n -th patient data $\mathbf{X}_n = ((\mathbf{X}_n^p, \mathbf{X}_n^g), c_n, t^n)$, as described in Section 3.1.1 and 3.1.2. After integrating H_n^p and H_n^g into the multimodal features H_n by concatenation, we estimate the hazard function $h_n(t|H_n) = h_n(T = t|T \geq t, H_n) \in [0, 1]$ which is viewed as the probability of death event occurring at around time point t . Instead of predicting the overall survival time t^n , the survival prediction task is to estimate the ordinal risk of death via the cumulative survival function:

$$S_n(t|H_n) = \prod_{z=1}^t (1 - h_n(z|H_n)) \quad (3)$$

3.2. Optimal Transport-based Co-Attention

To identify TME-related instances of WSIs, we utilize the global potential structure consistency between TME interactions of pathology and genes co-expressions as the evidence of selecting instances, in which optimal transport is used to learn the optimal matching flow between a WSI feature bag $\mathbf{B}_n^p \in \mathbb{R}^{M_p \times d}$ and a genomic bag $\mathbf{B}_n^g \in \mathbb{R}^{M_g \times d}$. Formally, an optimal transport from the WSI feature bag $\mathbf{B}_n^p = [\mathbf{b}_{n,1}^p, \mathbf{b}_{n,2}^p, \dots, \mathbf{b}_{n,M_p}^p]$ to the genomic embedding bag $\mathbf{B}_n^g = [\mathbf{b}_{n,1}^g, \mathbf{b}_{n,2}^g, \dots, \mathbf{b}_{n,M_g}^g]$ can be defined by the

discrete Kantorovich formulation [21] to search the overall optimal matching flow \mathbf{P}_n between \mathbf{B}_n^p and \mathbf{B}_n^g :

$$\mathcal{W}(\mathbf{B}_n^p, \mathbf{B}_n^g) = \min_{\mathbf{P}_n \in \Pi(\mu_p, \mu_g)} \langle \mathbf{P}_n, \mathbf{C}_n \rangle_F \quad (4)$$

where $\mathbf{C}_n \geq 0 \in \mathbb{R}^{M_p \times M_g}$ is a cost matrix given by $\mathbf{C}_n^{u,v} = c(\mathbf{b}_{n,u}^p, \mathbf{b}_{n,v}^g)$ with a ground distance metric $c(\cdot)$, such as l_2 -distance, that measures the distance of local pairwise instances $\mathbf{b}_{n,u}^p \in \mathbf{B}_n^p$ in the WSI bag and the genomic one of a unique functional category $\mathbf{b}_{n,v}^g \in \mathbf{B}_n^g$.

Here $\Pi(\mu_p, \mu_g) = \{\mathbf{P}_n \in \mathbb{R}_+^{M_p \times M_g} | \mathbf{P}_n \mathbf{1}_{M_g} = \mu_p, \mathbf{P}_n^\top \mathbf{1}_{M_p} = \mu_g\}$ involves the marginal constraints of total mass equality between marginal distributions, μ_p and μ_g , for the WSI bag and the genomic bag, and $\mathbf{1}_k$ is a k -dimensional vector of ones. Specifically, $\Pi(\mu_p, \mu_g)$ refers to the set of joint probabilistic couplings between the two marginal empirical distributions of pathology data and genomic data. Intuitively, it describes how to distribute instances of a WSI bag \mathbf{B}_n^p to the genomic embedding of a genomic bag \mathbf{B}_n^g based on the cost matrix \mathbf{C}_n , under the marginal constraints of total mass equality between distributions μ_p and μ_g . Note that $\langle \cdot \rangle_F$ refers to the Frobenius dot product, and thus Eq. 4 encourages to achieve the overall minimum matching cost by finding optimal matching flow based on local pairwise similarity. In this way, the optimal transport of Eq. 4 is able to enforce a trade-off of instances within modality, which allows the model to capture the potential structure of visual interactions for histology and co-expressions for genomic data.

Once acquiring the optimal matching flow \mathbf{P}_n^* , informative instances of a WSI are identified by $\hat{\mathbf{B}}_n^p = \mathbf{P}_n^\top \mathbf{B}_n^p$ to represent the WSI, which also aligns pathology distribution to genomics distribution under preserving the potential structure across modalities for alleviating heterogeneity.

As such, the proposed OT-based Co-Attention simultaneously addresses two issues: not only does it achieve the effective representation for gigapixel WSIs by identifying informative instances via OT matching, but marginal constraints of OT enable the selected instances to have global structure consistency with genomic data.

3.3. Optimization over Micro-Batch

Due to the large size of the WSI bag, it is difficult to apply optimal transport for matching histology data and genomic data. Recent work [48] has validated that MIL on WSI can benefit from multiple subsets of a bag randomly sampled from the original WSI bag. It provides a way to train the model over Micro-Batch of a WSI, where micro-batch is defined as a subset sampled from a bag of WSI instances. Furthermore, inspired by a variant [13] of OT that offers a solution to approximate the original OT over mini-batches with a theoretical guarantee of convergence, we propose to use the variant UMBOT formulation [13] over micro-batch of WSI instead of the original OT over

all instances of a WSI. Then the Eq. 4 becomes:

$$\begin{aligned} \mathcal{W}^m(B_n^{p,m}, \mathbf{B}_n^g) = & \min_{\mathbf{P}_n^m \in \Pi(\mu_p^m, \mu_g)} \langle \mathbf{P}_n^m, \mathbf{C}_n^m \rangle_F \\ & + \epsilon KL(\mathbf{P}_n^m | \mu_p^m \otimes \mu_g) + \tau (D_\phi(\mathbf{P}_{n,p}^m | \mu_p^m) + D_\phi(\mathbf{P}_{n,g}^m | \mu_g)) \end{aligned} \quad (5)$$

where $B_n^{p,m}$ is micro-batch of size m sampled from \mathbf{B}_n^p . Similarly, $\mathbf{C}_n^m \in \mathbb{R}^{m \times M_g}$ is the cost matrix over micro-batch of WSI $B_n^{p,m}$ and \mathbf{B}_n^g , and the optimal matching flow $\mathbf{P}_n^m \in \mathbb{R}^{m \times M_g}$ is optimized over micro-batch as well. Here $\mathbf{P}_{n,p}^m$ and $\mathbf{P}_{n,g}^m$ represent the marginals of \mathbf{P}_n^m , and D_ϕ is Csiszàr divergences, τ and $\epsilon \geq 0$ are the coefficients of marginal penalization and entropic regularization, respectively. KL refers to Kullback-Leible divergence.

Computational complexity. With the benefit of UMBOT for approximating the original OT, the computational complexity of optimization is reduced from $\mathcal{O}(M^3 \log(M))$ to $\mathcal{O}(\frac{M}{m} \times m^2) = \mathcal{O}(M \times m)$, where $M = \max(M_p, M_g)$ and $m \ll M$. The whole optimization procedure of the proposed MOTCat is presented in Alg. 1

Loss function. Following the previous work [5], the overall loss function is formulated by NLL-loss [46]:

$$\begin{aligned} \mathcal{L} = & -\frac{m}{M_p} \sum_{n=1}^N \sum_{X_n^{p,m} \in \mathbf{X}_n^p} c_n \cdot \log(S_n^m(t_n | H_n^m)) \\ & -\frac{m}{M_p} \sum_{n=1}^N \sum_{X_n^{p,m} \in \mathbf{X}_n^p} \{(1 - c_n) \cdot \log(S_n^m(t_n - 1 | H_n^m)) \\ & + (1 - c_n) \cdot \log(h_n^m(t_n | H_n^m))\} \end{aligned} \quad (6)$$

where H_n^m refers to the multimodal features over micro-batch formulated in line 16 of Alg. 1. h_n^m and S_n^m are the hazard function and cumulative survival function over micro-batch, respectively.

4. Experiment

In this section, we first present datasets and settings following previous experimental protocols [5, 6] for fair comparisons. Next, we demonstrate the performance results of the proposed method compared with the state-of-the-art (SOTA) methods, including unimodal and multimodal approaches. After that, we investigate the effectiveness of each component in our method and the effect on the size of micro-batch strategy. Finally, from a statistical point of view, we present Kaplan-Meier survival curves and Logrank test to show the performance of survival analysis. Interpretable visualization for histology data and genomic data are presented in supplementary materials.

4.1. Datasets and Settings

Datasets. To demonstrate the performance of the proposed method, we conducted various experiments over five public cancer datasets from The Cancer Genome Atlas (TCGA)

that contains paired diagnostic WSIs and genomic data with ground-truth survival outcome: Bladder Urothelial Carcinoma (BLCA), Breast Invasive Carcinoma (BRCA), Uterine Corpus Endometrial Carcinoma (UCEC), Glioblastoma & Lower Grade Glioma (GBMLGG), and Lung Adenocarcinoma (LUAD). The number of cases for each cancer type is shown by N in Tab. 1. Note that cases of these cancer datasets used for all compared methods are not less than cases used in the proposed method. For genomic data, the number of unique functional categories M_g is set as 6 following [28, 5], including 1) Tumor Suppression, 2) Oncogenesis, 3) Protein Kinases, 4) Cellular Differentiation, 5) Transcription, and 6) Cytokines and Growth.

Evaluation. For each cancer dataset, we perform 5-fold cross-validation with a 4:1 ratio of training-validation sets and report the cross-validated concordance index (C-Index) and its standard deviation (std) to quantify the performance of correctly ranking the predicted patient risk scores with respect to overall survival.

Implementation. For each WSI, we first apply the OTSU’s threshold method to segment tissue regions, and then non-overlapping 256×256 patches are extracted from the tissue region over $20 \times$ magnification. Then, we use ImageNet-pretrained [10] ResNet-50 [17] and a 256-d fully-connected layer as the feature encoder f_p to extract the 1024-d embedding for each patch, where parameters of ResNet-50 are frozen. For genomic data, we adopt SNN [22] as the feature encoder f_g following the setting of [5].

During training, we follow the setting of [5] for a fair comparison. Specifically, we adopt Adam optimizer with the initial learning rate of 2×10^{-4} and weight decay of 1×10^{-5} . Due to the large size of WSIs, the batch size for WSIs is 1 with 32 gradient accumulation steps, and all experiments are trained for 20 epoches. The size of Micro-Batch m is set as 256. Regarding the hyper-parameters of OT in Eq. 5, the coefficient of marginal penalization τ is 0.5 for all cancer datasets, and the coefficient of entropic regularization ϵ is 0.05 for BLCA and LUAD, as well as 0.1 for BRCA, UCEC and GBMLGG.

4.2. Results

We compare our method against the unimodal baselines and the multimodal SOTA methods as follows:

Unimodal Baseline. For genomic data, we adopt SNN [22] that has been used previously for survival outcome prediction in the TCGA [4, 5], and SNNTrans [22, 36] which incorporates SNN as the feature extractor and TransMIL [36] as the global aggregation model for MIL. For histology, we compare the SOTA MIL methods AttnMIL [19], DeepAttnMISL [45], CLAM [32], TransMIL [36], and DTFD-MIL [48].

Multimodal SOTA. We compare four SOTA methods for multimodal survival outcome prediction including

Table 1. C-Index (mean \pm std) performance over five cancer datasets. Patho. and Geno. refer to pathology modality and genomic modality, respectively. The best results and the second-best results are highlighted in **bold** and in underline, respectively.

Model	Patho.	Geno.	BLCA ($N = 373$)	BRCA ($N = 956$)	UCEC ($N = 480$)	GBMLGG ($N = 569$)	LUAD ($N = 453$)
SNN* [22]		✓	0.618 \pm 0.022	0.624 \pm 0.060	<u>0.679 \pm 0.040</u>	0.834 \pm 0.012	0.611 \pm 0.047
SNNTrans* [22, 36]		✓	0.659 \pm 0.032	0.647 \pm 0.063	0.656 \pm 0.038	0.839 \pm 0.014	0.638 \pm 0.022
AttnMIL* [19]	✓		0.599 \pm 0.048	0.609 \pm 0.065	0.658 \pm 0.036	0.818 \pm 0.025	0.620 \pm 0.061
DeepAttnMISL [45]	✓		0.504 \pm 0.042	0.524 \pm 0.043	0.597 \pm 0.059	0.734 \pm 0.029	0.548 \pm 0.050
CLAM-SB* [32]	✓		0.559 \pm 0.034	0.573 \pm 0.044	0.644 \pm 0.061	0.779 \pm 0.031	0.594 \pm 0.063
CLAM-MB* [32]	✓		0.565 \pm 0.027	0.578 \pm 0.032	0.609 \pm 0.082	0.776 \pm 0.034	0.582 \pm 0.072
TransMIL* [36]	✓		0.575 \pm 0.034	<u>0.666 \pm 0.029</u>	0.655 \pm 0.046	0.798 \pm 0.043	0.642 \pm 0.046
DTFD-MIL* [48]	✓		0.546 \pm 0.021	0.609 \pm 0.059	0.656 \pm 0.045	0.792 \pm 0.023	0.585 \pm 0.066
Pathomic [4]	✓	✓	0.586 \pm 0.062	-	-	0.826 \pm 0.009	0.543 \pm 0.065
PONET [34]	✓	✓	0.643 \pm 0.037	-	-	-	0.646 \pm 0.047
Porpoise* [6]	✓	✓	0.636 \pm 0.024	0.652 \pm 0.042	0.695 \pm 0.032	0.834 \pm 0.017	0.647 \pm 0.031
MCAT* [5]	✓	✓	<u>0.672 \pm 0.032</u>	0.659 \pm 0.031	0.649 \pm 0.043	<u>0.835 \pm 0.024</u>	<u>0.659 \pm 0.027</u>
MOTCat (Ours)	✓	✓	0.683 \pm 0.026	0.673 \pm 0.006	0.675 \pm 0.040	0.849 \pm 0.028	0.670 \pm 0.038

MCAT [5], Pathomic [4], Porpoise [6] and PONET [34].

Results are shown in Tab. 1, where the methods marked * are re-implemented and the best results are reported, and other results are obtained from their papers. Note that due to varying cancer datasets used in their works, not all results of the five datasets used in this work can be reported, and hence we use '-' to indicate the results are not reported in their papers.

Unimodal v.s. Multimodal. Compared with all unimodal approaches, the proposed method achieves the highest performance in 4 out of 5 cancer datasets, indicating the effective fusion of multimodal data in our method. Note that the overall performance of genomic data is better than that of histology, which validates the reasonability of using genomic data for guiding the selection of instances in WSIs. In particular, most multimodal methods are inferior to the unimodal model of genomics in UCEC dataset, suggesting the serious challenge in multimodal fusion. Nevertheless, the proposed method can achieve the comparative performance with genomic model on UCEC.

Multimodal SOTA v.s. MOTCat. In one-versus-all comparisons of multimodal models, the proposed method achieves superior performance on all benchmarks with 1.0%-2.6% performance gains except UCEC. In comparing the proposed method with the most similar work MCAT in multimodal fusion, our method gets better results on all datasets, indicating the effectiveness of learning the intra-modal potential structure consistency from a global perspective between TME-related interactions of histology and genomic co-expression.

4.3. Ablation Study

We first investigate the effectiveness of OT-based co-attention component (denoted by OT) and the micro-batch

strategy (denoted by MB). Furthermore, the effect on the size of micro-batch and the actual computational speed of our method are explored.

Component validation. As shown in Tab. 2, we observe that multimodal fusion can benefit from micro-batch strategy by comparing (a) and (b), and OT-based co-attention further improves the performance by approximating the original results with the averaged results over micro-batches via UMBOT. As a result, our method achieves the best overall performance. These improvements both demonstrate the effectiveness of each component for the proposed method.

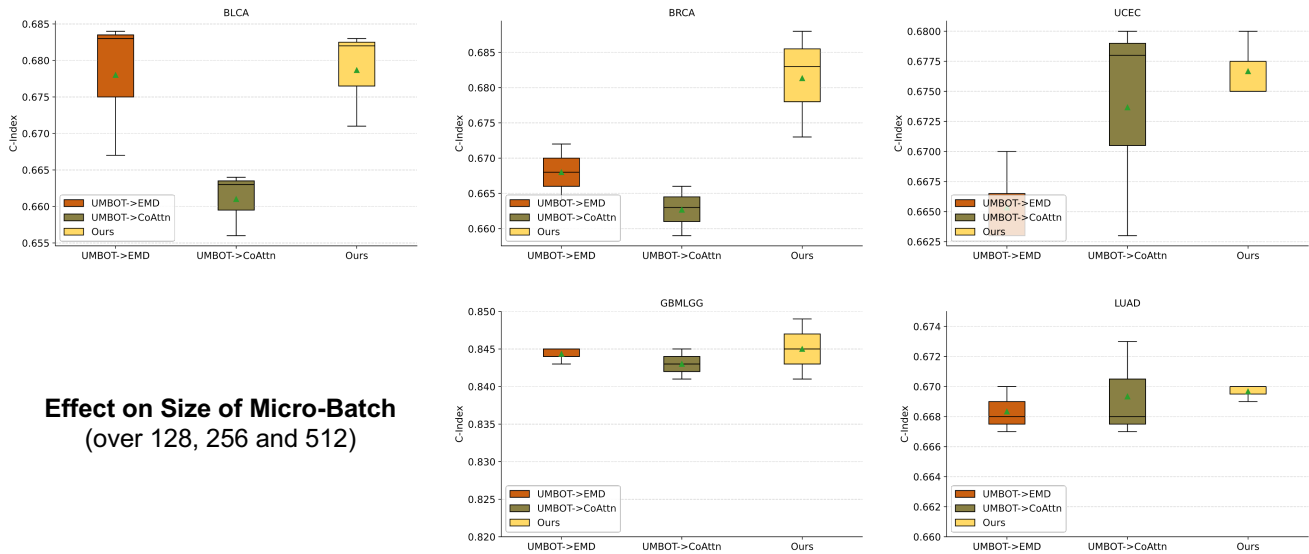
Additionally, recent work [48] has claimed that MIL on WSIs can benefit from learning over sub-bags in a WSI, which is also validated by performance increases in variant (b). The reason why it can profit from micro-batch might be that it increases the number of bags, so that more features can be used for survival analysis.

Size of Micro-batch. To show the robustness of approximation to the size of Micro-batch strategy, we compare several variants of the proposed method over the varying sizes of micro-batch, as shown in the quantitative analysis of Fig. 3. The qualitative analysis of the effect on size of Micro-Batch can be found in the supplementary materials, which also achieves the consistent conclusion.

For the quantitative analysis in Fig. 3, two variants of the proposed method are compared: 1) **UMBOT** \rightarrow **EMD**: To demonstrate the robustness of different OT variants, we replace UMBOT used in our method with the original OT (i.e. EMD) for comparison, which is equivalent to the Eq. (4) done over Micro-batch. 2) **UMBOT** \rightarrow **CoAttn**: To compare the robustness of the proposed OT-based co-attention and the original co-attention, we replace UMBOT with the co-attention used in MCAT [5] and train the model over Micro-batch for comparison.

Table 2. Ablation study assessing C-Index (mean \pm std) performance of MOTCat over five datasets, in which version (c) in bold is the proposed method with full components, and (a) and (b) refer to its variants. The best results are marked in **bold**.

Variants	OT	MB	BLCA	BRCA	UCEC	GBMLGG	LUAD
(a) MCAT			0.672 \pm 0.032	0.659 \pm 0.031	0.649 \pm 0.043	0.835 \pm 0.024	0.659 \pm 0.027
(b) MOTCat w/o OT		✓	0.663 \pm 0.025	0.666 \pm 0.017	0.663 \pm 0.031	0.845 \pm 0.023	0.673 \pm 0.036
(c) MOTCat	✓	✓	0.683 \pm 0.026	0.673 \pm 0.006	0.675 \pm 0.040	0.849 \pm 0.028	0.670 \pm 0.038



Effect on Size of Micro-Batch
(over 128, 256 and 512)

Figure 3. Boxplots of C-Index over various sizes of Micro-batch for our method and two variants on five cancer datasets, in which the green triangle means the averaged results of all sizes and the horizontal line in the box represents the median result.

From the results of Fig. 3, we can see that 1) our method achieves the best averaged performance of various sizes compared with the two variants, especially on BRCA with the most samples of 956. 2) Further, on UCEC and LUAD datasets, the proposed method gets the most robust results. 3) Although the second variant UMBOT \rightarrow CoAttn shows slightly better robustness on BLCA and BRCA, the performance of our method surpasses it by a large margin. In a nutshell, the proposed method achieves a better trade-off between performance and robustness, compared with different OT variants or the co-attention of MCAT.

Computational Speed. We measure the actual averaged computational speed for 10 WSIs (about 150k patches in total) on one GPU of NVIDIA GeForce RTX 3090. The proposed method is compared with a variant 'org-OT' that replaces UMBOT by the original OT. However, the original OT with the high complexity takes extreme long time for one WSI (about 20k patches), and thus we fail to measure the actual computational time for the original OT in a short time. Here we present the speed of our method, which suggests our method makes it practicable to apply OT to histology and genomics. Specifically, the training speed of the proposed method is 6540 p/s (p/s refers to the number

of processing patches per second) and the inference speed is 11885 p/s.

4.4. Statistical Analysis

To show a statistical difference in patient stratification performance, we visualize the Kaplan-Meier survival curves for different methods, in which patients are separated into two groups of low-risk and high-risk based on predicted risk scores, and then the statistics on ground-truth survival time are presented for each group in Fig. 4. The Logrank test is conducted to measure the statistically significant difference in two groups of patients, where a lower P-value indicates better performance of patient stratification. From the results of Fig. 4, intuitively our method separates patients of low and high risk more clearly on all datasets. In the Logrank test, our method achieves a lower P-value on all datasets in comparison to MCAT, especially on BLCA, BRCA and UCEC with a large margin of magnitude.

5. Conclusion

In this paper, we present a novel Multimodal Optimal Transport-based Co-Attention Transformer with global structure consistency to tackle two important issues in mul-

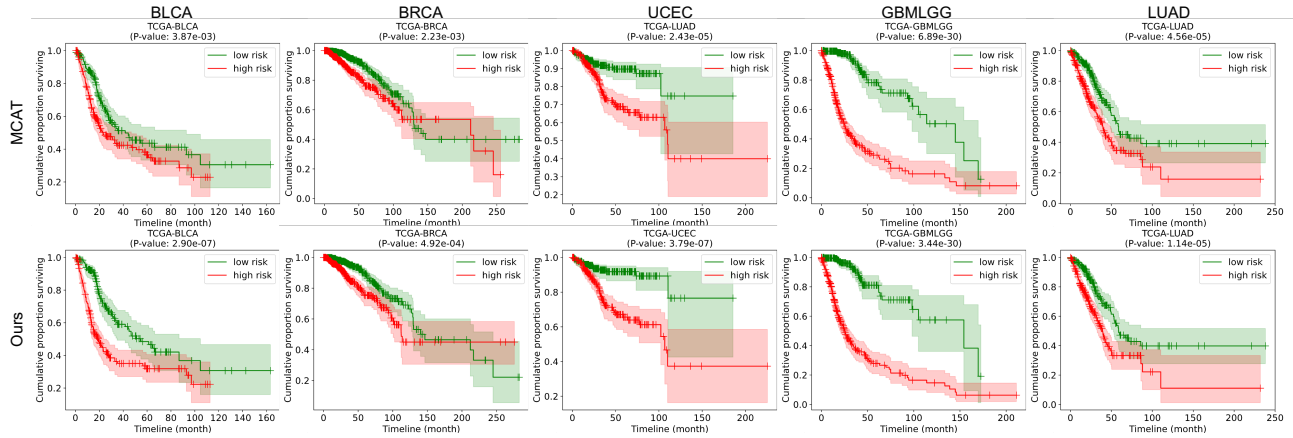


Figure 4. Kaplan-Meier Analysis on five cancer datasets, where patient stratifications of low risk (green) and high risk (red) are presented. Shaded areas refer to the confidence intervals. P-value < 0.05 means significant statistical difference in two groups, and lower P-value is better. (Zoom in for better viewing.)

timodal learning of histology and genomics for survival prediction. First, we utilize a new OT-based Co-Attention to match pairwise instances between histology and genomics for selecting informative instances strongly related to tumor microenvironment, which provides a way to effectively represent gigapixel WSIs. Second, optimal transport offers a global awareness for modeling potential structure within modality, i.e. pathological interactions and genomic co-expression. Furthermore, to enable applying OT-based Co-Attention in practice, we propose a robust and efficient implementation over micro-batch of WSI bags, in which the approximation over micro-batch might provide a solution to update the parameters of extractor for patch images in an end-to-end way instead of extracting feature embeddings offline, which will be further explored in the future.

Acknowledgement

This work was supported by National Natural Science Foundation of China (No. 62202403), Shenzhen Science and Technology Innovation Committee (Project No. SGDX20210823103201011) and Hong Kong Innovation and Technology Fund (No. PRP/034/22FX).

References

- [1] Khalid AbdulJabbar, Shan E Ahmed Raza, Rachel Rosenthal, Mariam Jamal-Hanjani, Selvaraju Veeriah, Ayse Akarca, Tom Lund, David A Moore, Roberto Salgado, Maise Al Bakir, et al. Geospatial immune variability illuminates differential evolution of lung adenocarcinoma. *Nature medicine*, 26(7):1054–1062, 2020. 1
- [2] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Mirafior, 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, 2019. 3

- [3] Zongsheng Cao, Qianqian Xu, Zhiyong Yang, Yuan He, Xiaochun Cao, and Qingming Huang. Otkge: Multi-modal knowledge graph embeddings via optimal transport. In *Advances in Neural Information Processing Systems*, 2022. 2
- [4] Richard J. Chen, Ming Y. Lu, Jingwen Wang, Drew F. K. 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*, 41(4):757–770, 2022. 1, 3, 6, 7
- [5] Richard J Chen, Ming Y Lu, Wei-Hung Weng, Tiffany Y Chen, Drew FK Williamson, Trevor Manz, Maha Shady, and Faisal Mahmood. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4025, 2021. 1, 2, 3, 5, 6, 7
- [6] Richard J Chen, Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Jana Lipkova, Zahra Noor, Muhammad Shaban, Maha Shady, Mane Williams, Bumjin Joo, et al. Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell*, 40(8):865–878, 2022. 1, 3, 6, 7
- [7] Philip Chikontwe, Meejeong Kim, Soo Jeong Nam, Heunjeong Go, and Sang Hyun Park. Multiple instance learning with center embeddings for histopathology classification. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part V 23*, pages 519–528. Springer, 2020. 3
- [8] Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of computation*, 87(314):2563–2609, 2018. 4
- [9] Cox R David et al. Regression models and life tables (with discussion). *Journal of the Royal Statistical Society*, 34(2):187–220, 1972. 3

- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [11] Jiali Duan, Liqun Chen, Son Tran, Jinyu Yang, Yi Xu, Belinda Zeng, and Trishul Chilimbi. Multi-modal alignment using representation codebook. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15651–15660, 2022. 2
- [12] David Faraggi and Richard Simon. A neural network model for survival data. *Statistics in medicine*, 14(1):73–82, 1995. 3
- [13] Kilian Fatras, Thibault Séjourné, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *International Conference on Machine Learning*, pages 3186–3197. PMLR, 2021. 2, 5
- [14] Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning with a wasserstein loss. *Advances in neural information processing systems*, 28, 2015. 4
- [15] Konrad Gadzicki, Razieh Khamsehashari, and Christoph Zetzsche. Early vs late fusion in multimodal convolutional neural networks. In *2020 IEEE 23rd international conference on information fusion (FUSION)*, pages 1–6. IEEE, 2020. 3
- [16] Noriaki Hashimoto, Daisuke Fukushima, Ryoichi Koga, Yusuke Takagi, Kaho Ko, Kei Kohno, Masato Nakaguro, Shigeo Nakamura, Hidekata Hontani, and Ichiro Takeuchi. Multi-scale domain-adversarial multiple-instance cnn for cancer subtype classification with unannotated histopathological images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3852–3861, 2020. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 630–645. Springer, 2016. 6
- [18] Shih-Cheng Huang, Anuj Pareek, Saeed Seyyedi, Imon Banerjee, and Matthew P Lungren. Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. *NPJ digital medicine*, 3(1):136, 2020. 3
- [19] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018. 2, 3, 6, 7
- [20] Sunghoon Joo, Eun Sook Ko, Soonhwan Kwon, Eunjoon Jeon, Hyungsik Jung, Ji-Yeon Kim, Myung Jin Chung, and Young-Hyuck Im. Multimodal deep learning models for the prediction of pathologic response to neoadjuvant chemotherapy in breast cancer. *Scientific reports*, 11(1):18800, 2021. 3
- [21] Leonid V Kantorovich. On the translocation of masses. *Journal of mathematical sciences*, 133(4):1381–1382, 2006. 5
- [22] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. *Advances in neural information processing systems*, 30, 2017. 6, 7
- [23] D. G. Kleinbaum and M. Klein. *Survival analysis : a self-learning text / 2nd ed.* Survival analysis : a self-learning text / 2nd ed, 2011. 3
- [24] Ashnil Kumar, Michael Fulham, Dagan Feng, and Jinman Kim. Co-learning feature fusion maps from pet-ct images of lung cancer. *IEEE Transactions on Medical Imaging*, 39(1):204–217, 2019. 3
- [25] Hajime Kuroda, Tsengelmaa Jamiyan, Rin Yamaguchi, Akinari Kakumoto, Akihito Abe, Oi Harada, and Atsuko Masunaga. Tumor-infiltrating b cells and t cells correlate with postoperative prognosis in triple-negative carcinoma of the breast. *BMC cancer*, 21(1):1–10, 2021. 1
- [26] Bin Li, Yin Li, and Kevin W Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2021. 2, 3
- [27] Ruiqing Li, Xingqi Wu, Ao Li, and Minghui Wang. Hfb-surv: hierarchical multimodal fusion with factorized bilinear models for cancer survival prediction. *Bioinformatics*, 38(9):2587–2594, 2022. 1, 3
- [28] Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P Mesirov, and Pablo Tamayo. The molecular signatures database hallmark gene set collection. *Cell systems*, 1(6):417–425, 2015. 2, 5, 6
- [29] Tsung-Chieh Lin, Yuan-Ming Yeh, Wen-Lang Fan, Yu-Chan Chang, Wei-Ming Lin, Tse-Yen Yang, and Michael Hsiao. Ghrelin upregulates oncogenic aurora a to promote renal cell carcinoma invasion. *Cancers*, 11(3):303, 2019. 2
- [30] Jana Lipkova, Richard J Chen, Bowen Chen, Ming Y Lu, Matteo Barbieri, Daniel Shao, Anurag J Vaidya, Chengkuan Chen, Luoting Zhuang, Drew FK Williamson, et al. Artificial intelligence for multimodal data integration in oncology. *Cancer Cell*, 40(10):1095–1110, 2022. 2, 3
- [31] Dan Liu, Xue Yang, and Xiongzi Wu. Tumor immune microenvironment characterization identifies prognosis and immunotherapy-related gene signatures in melanoma. *Frontiers in immunology*, 12:663495, 2021. 2
- [32] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature biomedical engineering*, 5(6):555–570, 2021. 2, 3, 6, 7
- [33] Yukiko Oya, Yoku Hayakawa, and Kazuhiko Koike. Tumor microenvironment in gastric cancers. *Cancer science*, 111(8):2696–2707, 2020. 1, 2
- [34] Lin Qiu, Aminollah Khormali, and Kai Liu. Deep biological pathway informed pathology-genomic multimodal survival prediction. *arXiv preprint arXiv:2301.02383*, 2023. 1, 3, 7
- [35] Dhanesh Ramachandram and Graham W Taylor. Deep multimodal learning: A survey on recent advances and trends. *IEEE signal processing magazine*, 34(6):96–108, 2017. 3
- [36] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image

- classification. *Advances in neural information processing systems*, 34:2136–2147, 2021. 3, 6, 7
- [37] Yash Sharma, Aman Shrivastava, Lubaina Ehsan, Christopher A Moskaluk, Sana Syed, and Donald Brown. Cluster-to-conquer: A framework for end-to-end multi-instance learning for whole slide image classification. In *Medical Imaging with Deep Learning*, pages 682–698. PMLR, 2021. 3
- [38] Qi-Min Wang, LI Lv, Ying Tang, LI Zhang, and Li-Fen Wang. 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*, 17(2):1732–1740, 2019. 2
- [39] Siwen Wang, Caizhen Feng, Di Dong, Hailin Li, Jing Zhou, Yingjiang Ye, Zaiyi Liu, Jie Tian, and Yi Wang. Preoperative computed tomography-guided disease-free survival prediction in gastric cancer: a multicenter radiomics study. *Medical Physics*, 47(10):4862–4871, 2020. 2
- [40] Xi Wang, Hao Chen, Caixia Gan, Huangjing Lin, Qi Dou, Efstratios Tsougenis, Qitao Huang, Muyan Cai, and Pheng-Ann Heng. Weakly supervised deep learning for whole slide lung cancer image analysis. *IEEE transactions on cybernetics*, 50(9):3950–3962, 2019. 3
- [41] Yun Wang, Tong Zhang, Xueya Zhang, Zhen Cui, Yuge Huang, Pengcheng Shen, Shaoxin Li, and Jian Yang. Wasserstein coupled graph learning for cross-modal retrieval. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1793–1802. IEEE, 2021. 2
- [42] Zongjie Wang, Sharif Ahmed, Mahmoud Labib, Hansen Wang, Xiyue Hu, Jiarun Wei, Yuxi Yao, Jason Moffat, Edward H Sargent, and Shana O Kelley. Efficient recovery of potent tumour-infiltrating lymphocytes through quantitative immunomagnetic cell sorting. *Nature Biomedical Engineering*, 6(2):108–117, 2022. 1
- [43] Zhiqin Wang, Ruiqing Li, Minghui Wang, and Ao Li. Gpbn: deep bilinear network integrating both genomic data and pathological images for breast cancer prognosis prediction. *Bioinformatics*, 37(18):2963–2970, 2021. 3
- [44] Gang Xu, Zhigang Song, Zhuo Sun, Calvin Ku, Zhe Yang, Cancheng Liu, Shuhao Wang, Jianpeng Ma, and Wei Xu. Camel: A weakly supervised learning framework for histopathology image segmentation. In *Proceedings of the IEEE/CVF International Conference on computer vision*, pages 10682–10691, 2019. 3
- [45] Jiawen Yao, Xinliang Zhu, Jitendra Jonnagaddala, Nicholas Hawkins, and Junzhou Huang. Whole slide images based cancer survival prediction using attention guided deep multiple instance learning networks. *Medical Image Analysis*, 65:101789, 2020. 6, 7
- [46] Shekoufeh Gorgi Zadeh and Matthias Schmid. Bias in cross-entropy-based training of deep survival networks. *IEEE transactions on pattern analysis and machine intelligence*, 43(9):3126–3137, 2020. 6
- [47] Yangyang Zeng, Yulan Zeng, Hang Yin, Fengxia Chen, Qingqing Wang, Xiaoyan Yu, and Yunfeng Zhou. Exploration of the immune cell infiltration-related gene signature in the prognosis of melanoma. *Aging (albany NY)*, 13(3):3459, 2021. 2
- [48] Hongrun Zhang, Yanda Meng, Yitian Zhao, Yihong Qiao, Xiaoyun Yang, Sarah E Coupland, and Yalin Zheng. Dtfidmil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18802–18812, 2022. 2, 3, 5, 6, 7
- [49] Jie Zhang and Kun Huang. Normalized imqcm: An algorithm for detecting weak quasi-cliques in weighted graph with applications in gene co-expression module discovery in cancers. *Cancer informatics*, 13:CIN–S14021, 2014. 2
- [50] Hanci Zheng, Zongying Lin, Qizheng Zhou, Xingchen Peng, Jianghong Xiao, Chen Zu, Zhengyang Jiao, and Yan Wang. Multi-transsp: Multimodal transformer for survival prediction of nasopharyngeal carcinoma patients. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VII*, pages 234–243. Springer, 2022. 2
- [51] Meng Zhou, Zicheng Zhang, Siqi Bao, Ping Hou, Congcong Yan, Jianzhong Su, and Jie Sun. Computational recognition of lncrna signature of tumor-infiltrating b lymphocytes with potential implications in prognosis and immunotherapy of bladder cancer. *Briefings in Bioinformatics*, 22(3):bbaa047, 2021. 2
- [52] Qi Zhu, Heyang Wang, Bingliang Xu, Zhiqiang Zhang, Wei Shao, and Daoqiang Zhang. Multimodal triplet attention network for brain disease diagnosis. *IEEE Transactions on Medical Imaging*, 41(12):3884–3894, 2022. 3
- [53] Xinliang Zhu, Jiawen Yao, and Junzhou Huang. Deep convolutional neural network for survival analysis with pathological images. In *2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 544–547. IEEE, 2016. 3
- [54] Xinliang Zhu, Jiawen Yao, Feiyun Zhu, and Junzhou Huang. Wsisa: Making survival prediction from whole slide histopathological images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7234–7242, 2017. 1, 2
- [55] Yingli Zuo, Yawen Wu, Zixiao Lu, Qi Zhu, Kun Huang, Daoqiang Zhang, and Wei Shao. Identify consistent imaging genomic biomarkers for characterizing the survival-associated interactions between tumor-infiltrating lymphocytes and tumors. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II*, pages 222–231. Springer, 2022. 1, 2, 3