

DRIM: Learning Disentangled Representations from Incomplete Multimodal Healthcare Data

Lucas Robinet^{1,2,3}, Ahmad Berjaoui², Ziad Kheil^{1,3}, and Elizabeth Cohen-Jonathan Moyal^{1,3}

¹ IUCT-Oncopole-Institut Claudius Regaud, Toulouse, France
`{robinet.lucas,moyal.elisabeth}@iuct-oncopole.fr`

² IRT Saint-Exupéry, Toulouse, France
`ahmad.berjaoui@irt-saintexupery.com`

³ INSERM UMR 1037, Cancer Research Center of Toulouse (CRCT), Toulouse, France
`ziad.kheil@inserm.fr`

Abstract. Real-life medical data is often multimodal and incomplete, fueling the growing need for advanced deep learning models capable of integrating them efficiently. The use of diverse modalities, including histopathology slides, MRI, and genetic data, offers unprecedented opportunities to improve prognosis prediction and to unveil new treatment pathways. Contrastive learning, widely used for deriving representations from paired data in multimodal tasks, assumes that different views contain the same task-relevant information and leverages only shared information. This assumption becomes restrictive when handling medical data since each modality also harbors specific knowledge relevant to downstream tasks. We introduce DRIM, a new multimodal method for capturing these shared and unique representations, despite data sparsity. More specifically, given a set of modalities, we aim to encode a representation for each one that can be divided into two components: one encapsulating patient-related information common across modalities and the other, encapsulating modality-specific details. This is achieved by increasing the shared information among different patient modalities while minimizing the overlap between shared and unique components within each modality. Our method outperforms state-of-the-art algorithms on glioma patients survival prediction tasks, while being robust to missing modalities. To promote reproducibility, the code is made publicly available at <https://github.com/Lucas-rbnt/DRIM>.

Keywords: Disentanglement · Deep Learning · Multimodal · Transformers · Glioma

1 Introduction

Cancer diagnosis and prognosis are usually determined by an ensemble of data obtained from various sources including MRI, histopathology slides, and molecular profiles. Combining these modalities allows to derive a panoply of biomark-

ers, offering a deeper understanding of the tumor landscape. Thereby, it enables clinicians to provide patients with the utmost support throughout their pathology [12]. In fields beyond medicine, there are several deep learning models and techniques capable of merging multimodal data [15,25,31,18]. However, in a medical context, the prevalent challenge of integrating incomplete heterogeneous modalities steers deep learning techniques towards simpler multimodal or even unimodal-based models [30,2]. While some multimodal approaches rely on late fusion [27] i.e. averaging outputs from each modality, most popular methods focus on merging modalities in the latent space [29,8,5,9,21]. Indeed, MultiSurv [29] explores the utility of straightforward fusion methods such as mean, sum, product, and maximum, testing their capacity to handle missing modalities within survival prediction tasks. Although simple, another fusion technique that has showcased its effectiveness in the medical domain involves the concatenation of the representations [27]. Finally, [31] introduced a parameterized approach called Tensor, which leverages the outer product of representations. This method was later applied in [9] to combine histology slides and genomic data. Nonetheless, this method falls short in naturally accommodating missing modalities and faces a critical limitation as the amount of parameters it requires increases exponentially with each added modality. This excessively high number of parameters hinders its effectiveness for combining multiple modalities as it drastically limits the representation dimensions.

Beyond mere fusion, leveraging auxiliary losses on representations to induce specific behaviors is a strategy adopted in previous works [8,5,32]. Hence, [8,32] embed a correlation function within representations to promote consistency. Conversely, [5] adopts a divergent strategy, denoted multimodal orthogonalization (MMO), striving for their representations to minimize redundancy and overlapping information across modalities.

In this paper, unlike to prior work, we make the natural assumption that information contained in a modality splits into patient-related aspects shared across modalities and unique features specific to each modality. Our key contributions are: (i) employing a pair of encoders for each modality —one shared, one unique. Optimizing the shared encoder can be viewed through the lens of Information Bottleneck [28], ensuring that representations are rich in patient-specific information yet minimally captures the modality’s characteristics. Concurrently, unique encoders aim to reduce mutual information between their representation and the corresponding shared one. (ii) Leveraging these disentangled representations, we introduce a dual-scale fusion approach: it initiates with the aggregation of shared information, later combined with unique information to craft a comprehensive representation pertinent for survival analysis. (iii) We employ an attention-based, scalable and parameterized fusion that intuitively models interactions between representations and natively manages missing modalities.

Gliomas, with their diverse prognosis biomarkers, range of grades, and complexity as described in the WHO 2021 classification [24], present a compelling use case for showcasing DRIM’s ability to effectively capture and disentangle intricate information.

2 Method




Fig. 1: An overview of the proposed method. (a) describes the learning of shared and unique components, while (b) and (c) propose a supervised and an unsupervised alternative to couple these representations to a specific task.

Consider a training minibatch of N samples and M modalities, with $x^m = \{x_i^m\}_{i=1}^N$ representing data from the N samples for a specific modality m . For each modality m , let $\phi^m : x^m \mapsto s^m \in \mathcal{R}^{N \times d}$ and $\xi^m : x^m \mapsto u^m \in \mathcal{R}^{N \times d}$ denote the shared and unique encoders respectively, with s_i^m and u_i^m , being the representations of the m -th modality for the i -th sample in the minibatch. DRIM makes it possible to learn disentangled representations while being flexible to a range of tasks. Indeed, the DRIM-Surv approach (Fig. 1b) operates by fusing representations at two distinct scales to perform supervised end-to-end survival training. Alternatively, DRIM-U provides a wholly unsupervised option for learning meaningful shared and unique representations as illustrated in Fig. 1c. In both scenarios, the DRIM loss can be dissected into three distinct components as illustrated in Eq. 1: a task-related term, a shared term, and a unique term.

$$\mathcal{L}_{DRIM} = \mathcal{L}_T + \mathcal{L}_{sh} + \gamma \mathcal{L}_u \quad (1)$$

2.1 Shared Loss

To maximize mutual information between shared representations s^m , traditional methods optimize the InfoNCE objective [25] for each modality pair [15,8], lead-

ing to a computational complexity of $O(M^2)$. Thereby, in Eq. 2, we propose a new well-suited shared cost function (see Fig. 1a). Inspired by supervised contrastive learning [22] wherein class labels facilitate generalization to multiple positive pairs, we stipulate that each representation s^m finds positive pairs among representations stemming from distinct modalities yet originating from the same patient. Thus, all shared representations are concatenated to create the matrix $S = [s_1^1, s_2^1, \dots, s_1^2, s_2^2, \dots, s_{N-1}^M, s_N^M]^T \in \mathcal{R}^{N \cdot M \times d}$. Let $J := \{1, \dots, N \cdot M\}$ represent the indices of representations, $B(j) = \{b \in J : b \neq j, \kappa(b) = \kappa(j)\}$ define the set of its positive pairs, where $\kappa(j)$ maps j to the patient it belongs to, $\tau \in \mathcal{R}^+$ is a scalar temperature parameter. The shared loss term is formulated as follows:

$$\mathcal{L}_{sh} = \frac{1}{N \cdot M} \sum_{j \in J} \frac{-1}{M-1} \sum_{b \in B(j)} \log \frac{\exp(s_j^T s_b / \tau)}{\sum_{k \in J, k \neq j} \exp(s_j^T s_k / \tau)} \quad (2)$$

2.2 Unique Loss

Until now, we have successfully derived shared representations $\{s^m\}_{m=1}^M$. The unique representation u^m , extracted by ξ^m , parameterized by ψ^m , should highlight features specific to its modality m without encompassing the information contained in s^m . To this end, we use an adversarial objective [16], akin to the approach outlined in [26], to minimize the mutual information $I(s^m, u^m)$. A discriminator D^m , parameterized by θ_m is trained for each modality m to differentiate representations sampled from the joint distribution from those sampled from the product of the marginals as shown in Fig. 1a. Simulating samples from the product of the marginals $\mathbb{P}_S \mathbb{P}_U$ involves shuffling the batch of shared representations before pairing with unique ones as in [4]. This adversarial objective outlined in Eq. 3 drives the minimization of the Jensen-Shannon divergence $D_{JS}(\mathbb{P}_{S,U} \parallel \mathbb{P}_S \mathbb{P}_U)$, as evidenced in [16].

$$\begin{aligned} \mathcal{L}_u(\theta, \psi) = & - \sum_{m=1}^M \left[\mathbb{E}_{\mathbb{P}_{S,U}} [\log(1 - D^m(s^m, u^m)); \theta^m] \right. \\ & \left. + \mathbb{E}_{\mathbb{P}_S \mathbb{P}_U} [\log D^m(s^m, u^m); \theta^m] + \mathbb{E}_{\mathbb{P}_{S,U}} [\log D^m(s^m, u^m); \psi^m] \right] \quad (3) \end{aligned}$$

2.3 Task Loss

Merely separating shared and unique components does not guarantee the relevance of the learned unique representations. The training of ξ^m might still collapse, risking a meaningless latent space. Hence, to ensure we capture of pertinent information, each unique encoder is tied to a specific task. Indeed, in DRIM-Surv, unique encoders are also linked to a survival task, compelling them to identify modality-specific relevant features. To delve deeper, let $\{x_i, t_i, \delta_i\}_{i=1}^N$ denote our minibatch where t_i represents the time of last follow-up and δ_i indicates if the event was observed ($\delta_i = 1$) or censored ($\delta_i = 0$). Time is divided into P equidistant fixed intervals, and a neural network $f(x) = [\tilde{h}_1, \dots, \tilde{h}_P] \in [0, 1]^P$

aims to estimate the conditional hazard probability for each time interval. Let also $v : \mathcal{R} \rightarrow \{1, \dots, P\}$ be a function that maps a follow-up t to the interval in which it lies. The survival loss function used is an extension of [13] provided by [23] allowing predicted hazards to be non-proportional.

$$\mathcal{L}_T = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^{v(t_i)} \mathbb{1}_{\{v(t_i)=k, \delta_i=1\}} \log \tilde{h}_k + (1 - \mathbb{1}_{\{v(t_i)=k, \delta_i=1\}}) \log (1 - \tilde{h}_k) \quad (4)$$

Moreover, with the DRIM-U alternative, each ξ^m is also optimized to encode and decode inputs through a specific decoder $R^m(\xi^m(x^m)) = \hat{x}^m$, minimizing reconstruction error $\mathcal{L}_T = \sum_{m=1}^M \mathcal{L}_r^m$ and $\mathcal{L}_r^m = \sum_i \|\hat{x}_i^m - x_i^m\|_2^2$ as in Fig. 1c.

2.4 Multimodal Fusion

In DRIM-Surv, we merge representations on two scales: first, h_s fuses shared representations into a global shared vector $s \in \mathcal{R}^{N \times d}$; then, this is combined with unique representations via h_u to form a unified patient embedding, depicted in Fig. 1b. In both fusion, a masked transformer processes each representation as a sequence token, with a classification token $\langle \text{cls} \rangle$ added at the beginning, akin to a ViT approach [11]. The output of the transformer for this token acts as the fused representation, as shown in Fig. 1b. Specifically, h_u aggregates the M unique representations plus the global shared one s through $h_u([u^1, \dots, u^M, s]) = \text{softmax}(\frac{Q_{cls} K^T}{\sqrt{D}}) V$. where $Q_{cls} = q_{cls} W_Q$, $K = C W_K$ and $V = C W_V$ are linear transformations of the learnable query token $\langle \text{cls} \rangle$, and representations $C = [\langle \text{cls} \rangle, u^1, \dots, u^M, s]$ respectively and D the projection dimension. In practice, we use multiple attention heads. This fusion, termed MAFusion, offers scalability, computational efficiency, robust handling of missing modalities through masking and provides interpretability via attention maps.

3 Experiments

This study leverages multimodal glioma data from TCGA [7], specifically GBM and LGG collections, across four modalities: DNA methylation, (DNAm), RNA sequencing (RNA), histopathological slides (WSI) and MRI. MRI data was completed from the BraTS dataset [3], selecting only TCGA-listed patients. Focusing on patient with at least two modalities, we kept 881 samples, finding coverage rates of 95% for DNAm, 89% for WSI, 75% for RNA, and 18% for MRI.

Genomics (DNAm and RNA): DNA methylation data were processed into 25,978 Beta values using the same method as in [29]. RNA data were processed to retain protein-coding genes with variance over 0.1, resulting in 16,304 FPKM-UQ normalized genes, then log-transformed and z-score normalized. Both genomics encoders consist of six 1D convolutional layers, with channel sizes of 8, 32, 64, 128, 256, and the final representation size. Each layer applies GELU activation [20], batch normalization, and dropout, ending with Adaptive Average Pooling to produce a fixed-size embedding.

WSI: Each patient was linked to a single, hand-selected image with minimal artifacts (pen marks, blurred zones, etc.). Chosen slides were subsampled, and a mask to isolate tissue zones was extracted using an OTSU filter. From these, 1000 high-resolution patches of size 256×256 were randomly selected, of which 100 were chosen based on a custom HSV intensity score (Fig. S1). A SimCLR technique [10] was employed to train a ResNet34 [19] on the chosen patches. Then, the WSI encoder, a single-layer transformer with 8 attention heads and no positional encoding, projects 10 randomly selected patch representations into a meaningful embedding.

MRI scans: Segmentation masks were used to crop a $64 \times 64 \times 64$ voxel region surrounding the tumor in Gd-T1w and FLAIR MRI scans. A SimCLR approach [10], similar to the WSI one, is used to train a MONAI ResNet10-3D [6] across the BraTS dataset [3]. In our DRIM experiments, we froze the encoder weights and added a 256-unit linear layer with dropout and GELU activation [20], followed by a dense layer.

Implementation details: Models were trained using an AdamW(lr= $1e^{-3}$, wd= $1e^{-2}$) optimizer over 30 epochs with a Cosine Annealing Scheduler and a batch size of 24. For all experiments, the representation size was set to 128, except in those involving Tensor [31], where it was reduced to 32 due to memory constraints—approaching 34,500 million parameters at a dimension of 128 (Table S1). The final layer of each model consisted of 20 units with sigmoid activation. The MAFusion utilized 16 attention heads of size 64 and a dense layer of size 128. In DRIM experiments, discriminators were also optimized via AdamW (lr= $1e^{-3}$, wd= $3e^{-4}$), and the loss coefficient γ is set empirically to 0.8. For the DRIM-U evaluation, encoders were pre-trained for 30 epochs as detailed in Section 2 and then frozen. Subsequently, a DRIM-Surv module was fine-tuned on top of these representations over 10 epochs for survival prediction.

Evaluation: First, the dataset is divided into 80% for training and 20% for testing. A five-fold cross-validation is then performed on the training set, yielding five models. Then, each of these trained models is evaluated on the hold-out test set. Performance metrics are reported as mean $\pm$ standard deviation. Evaluation metrics include the Concordance Index (C-index) [1], Integrated Brier Score (IBS) [14,17], and Integrated Negative Binomial Log-likelihood (INBLL) [17,23].

4 Results and discussions

4.1 Survival performance

The data in Table 1 demonstrate that the DRIM-Surv approach outperforms several other methods across all three assessed metrics. This superiority highlights the strong ability of our method to accurately discriminate patient outcomes (C-index) as well as to reliably forecast survival probabilities over time

Table 1: Comparison of mean performance metrics among survival models.

Method	Fusion	C-index ($\uparrow$)	Integrated BS ($\downarrow$)	INBLL ($\downarrow$)
Vanilla	Mean [29]	0.763 ± 0.010	0.110 ± 0.005	0.352 ± 0.026
	MMean [8]	0.757 ± 0.008	0.110 ± 0.006	0.361 ± 0.030
	Conc. [27]	0.760 ± 0.006	0.095 ± 0.005	0.319 ± 0.019
	Sum [29]	0.755 ± 0.007	0.097 ± 0.004	0.328 ± 0.030
	Max [29]	0.766 ± 0.006	0.097 ± 0.006	0.313 ± 0.026
	Tensor [9]	0.765 ± 0.006	0.093 ± 0.005	0.292 ± 0.015
	MAFusion	0.728 ± 0.009	0.100 ± 0.005	0.382 ± 0.029
MMO Loss [5]	Mean [29]	0.759 ± 0.008	0.111 ± 0.005	0.358 ± 0.029
	Conc. [27]	0.764 ± 0.007	0.093 ± 0.004	0.311 ± 0.018
	Tensor [9]	0.764 ± 0.004	0.090 ± 0.007	0.286 ± 0.026
	MAFusion	0.728 ± 0.011	0.100 ± 0.008	0.389 ± 0.052
Sh. Loss only	MAFusion	0.750 ± 0.016	0.090 ± 0.006	0.324 ± 0.022
DRIM-U (ours)	MAFusion	0.763 ± 0.006	0.091 ± 0.06	0.286 ± 0.021
DRIM-Surv (ours)	MAFusion	0.774 ± 0.006	0.086 ± 0.004	0.285 ± 0.016

(IBS, INBLL). Notably, DRIM-Surv not only slightly exceeds the performance of Tensor [31] in all metrics, but it also offers advantages in terms of scalability to additional modalities, reduced training time, and a significantly lower parameter count (Table S1 and Table S2). Indeed, except for the new encoders parameters, incorporating an additional modality into the DRIM framework leaves its parameter count unchanged, while the parameter total for Tensor [31] would escalate from 38 million—already near 70 times larger than those of our MAFusion—to exceed one billion (Table S1). Additionally, despite using a representation of size 32, which is four times smaller than our method, TensorFusion [31] requires nearly twice the GPU memory (Table S2). The sensitivity analysis on the γ coefficient illustrates that, beyond the neural network architecture, our method effectively learns to disentangle the representations of each modality, thereby enhancing the reliability of the final prognosis (Table S3). Finally, our unsupervised DRIM-U approach also gives convincing results simply by being fine-tuned on a few epochs, validating the flexibility of our method.

4.2 Stratifying high-risk vs low-risk patients

Additionally, we evaluate the proficiency of our model to stratify patients into high-risk and low-risk categories, a crucial step towards advancing personalized medicine. Concretely, for a given model, patients are divided into two groups based on their cumulative predicted hazards: those in the top half are classified as high-risk, while those in the bottom half are considered low-risk. Survival curves for each group are derived using Kaplan-Meier estimators, with a log-rank test evaluating their differences and stratification efficacy. For DRIM-Surv, The logrank test yields a p-value of $2.9e^{-23}$, against $4.2e^{-22}$ for the Tensor [31]

w/ MMO [5] method, emphasizing the ability of our model to distinguish survival outcomes between the two risk groups (Fig. S2).

4.3 Robustness across varied input modality combinations

One of the primary benefits of our approach lies in its ability to smoothly manage missing modalities during inference, avoiding the use of zero-filled tensors required by other methods [31,5,27]. Indeed, DRIM goes beyond simple arranged modality combination; it adeptly models their interactions and efficiently extracts crucial information, showcasing advanced data understanding. Table 2

Table 2: Analysis of CS ($\frac{C^{index} + (1-IBS)}{2}$) metrics [27] on test samples containing at least all specified modalities, setting any extra modalities to zero for inference. The table displays only the top three competing models. For each combination of inputs, the percentage of occurrence in the training dataset is given (% train).

Modalities (% train)	Max	Conc. w/ MMO	Tensor w/ MMO	DRIM-Surv (ours)
WSI,MRI (2%)	0.702 ± 0.019	0.643 ± 0.039	0.619 ± 0.060	0.796 ± 0.016
WSI,RNA (2%)	0.795 ± 0.011	0.816 ± 0.010	0.811 ± 0.021	0.829 ± 0.006
DNA,m,WSI (17%)	0.798 ± 0.004	0.810 ± 0.007	0.804 ± 0.010	0.832 ± 0.008
DNA,m,RNA (9%)	0.843 ± 0.006	0.848 ± 0.009	0.860 ± 0.009	0.858 ± 0.011
DNA,m,MRI (2%)	0.740 ± 0.018	0.740 ± 0.014	0.740 ± 0.015	0.794 ± 0.025
RNA,MRI (0%)	0.783 ± 0.039	0.803 ± 0.041	0.755 ± 0.033	0.874 ± 0.024
DNA,m,WSI,RNA(55%)	0.838 ± 0.007	0.839 ± 0.010	0.853 ± 0.009	0.850 ± 0.015
DNA,m,WSI,MRI (3%)	0.767 ± 0.031	0.779 ± 0.022	0.761 ± 0.029	0.838 ± 0.020
DNA,m,RNA,MRI(1%)	0.824 ± 0.033	0.870 ± 0.036	0.823 ± 0.047	0.874 ± 0.032
WSI,RNA,MRI (1%)	0.828 ± 0.019	0.830 ± 0.027	0.771 ± 0.043	0.892 ± 0.026
All 4 modalities (8%)	0.886 ± 0.014	0.885 ± 0.017	0.828 ± 0.050	0.904 ± 0.029

reveals that while other models show comparable performance on their most frequent training combination (55%), DRIM uniquely excels across all input combinations, surpassing its competitors. For instance, when inferring on WSI and MRI data, DRIM attains a CS score of 0.796 ± 0.016 , far surpassing the 0.619 ± 0.060 achieved by Tensor [31] w/ MMO [5]. Furthermore, our method is the only one that almost systematically improves its performance when a modality is added to a given combination. This distinction demonstrates that, whereas other models rely heavily on *learning* from a fixed and ordered sequence of modalities, DRIM-Surv stands out in its capability to specifically extract and combine information from diverse inputs.

5 Conclusions and perspectives

We propose DRIM, an approach for disentangling data representations from highly heterogeneous and incomplete sources. Our method not only enhances

the forecasting of patient outcomes through an adept handling of input modalities but also demonstrates improved scalability and flexibility with the addition of new modalities, unlike classic multimodal techniques. This versatility ensures its applicability to a broad spectrum of datasets featuring diverse modalities and tasks. Finally, we aspire that the ongoing efforts to disentangle learned representations of modalities will pave the way for deeper analysis and enhanced interpretability. This development promises to unlock precise, individualized patient assessments, thereby bridging the gap to personalized medicine and practical clinical adoption.

Disclosure of Interests. Elizabeth Cohen-Jonathan Moyal served as a member of expert board for Novocure, received lecture fees from Accuray and Novocure, travel fees from Novocure and research grants from Astra Zeneca, Novocure, Bayer and Incyte . She also received research grants from ARC foundation. All other authors have no competing interests to declare.

References

1. Antolini, L., Boracchi, P., Biganzoli, E.: A time-dependent discrimination index for survival data. *Statistics in Medicine* **24**(24) (2005)
2. Bae, S., Choi, Y.S., Ahn, S.S., Chang, J.H., Kang, S.G., Kim, E.H., Kim, S.H., Lee, S.K.: Radiomic MRI Phenotyping of Glioblastoma: Improving Survival Prediction. *Radiology* **289**(3), 797–806 (Dec 2018)
3. Baid, U.e.a.: The RSNA-ASNR-MICCAI BraTS 2021 Benchmark on Brain Tumor Segmentation and Radiogenomic Classification (Sep 2021), arXiv:2107.02314 [cs]
4. Brakel, P., Bengio, Y.: Learning Independent Features with Adversarial Nets for Non-linear ICA (Oct 2017), arXiv:1710.05050 [stat]
5. Braman, N., Gordon, J.W.H., Goossens, E.T., Willis, C., Stumpe, M.C., Venkataraman, J.: Deep Orthogonal Fusion: Multimodal Prognostic Biomarker Discovery Integrating Radiology, Pathology, Genomic, and Clinical Data. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 667–677 (2021)
6. Cardoso, M.e.a.: MONAI: An open-source framework for deep learning in healthcare (Nov 2022), arXiv:2211.02701 [cs]
7. Chang, K., The Cancer Genome Atlas Research Network, Genome Characterization Center, Genome Data Analysis Center, Sequencing Center, Data Coordinating Center, Tissue Source Site, Biospecimen Core Resource Center: The Cancer Genome Atlas Pan-Cancer analysis project. *Nature Genetics* **45**(10), 1113–1120 (Oct 2013)
8. Cheerla, A., Gevaert, O.: Deep learning with multimodal representation for pan-cancer prognosis prediction. *Bioinformatics* **35**(14), i446–i454 (Jul 2019)
9. Chen, R.J., Lu, M.Y., Wang, J., Williamson, D.F.K., Rodig, S.J., Lindeman, N.I., Mahmood, F.: Pathomic Fusion: An Integrated Framework for Fusing Histopathology and Genomic Features for Cancer Diagnosis and Prognosis (Sep 2020), arXiv:1912.08937 [cs, q-bio]
10. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A Simple Framework for Contrastive Learning of Visual Representations. In: *Proceedings of the 37th International Conference on Machine Learning*. pp. 1597–1607. PMLR (Nov 2020), iSSN: 2640-3498

11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale (Oct 2020)
12. Gallego, O.: Nonsurgical treatment of recurrent glioblastoma. *Current Oncology* (Toronto, Ont.) **22**(4), e273–281 (Aug 2015). <https://doi.org/10.3747/co.22.2436>
13. Gensheimer, M.F., Narasimhan, B.: A scalable discrete-time survival model for neural networks. *PeerJ* **7**, e6257 (Jan 2019)
14. Gerds, T.A., Schumacher, M.: Consistent Estimation of the Expected Brier Score in General Survival Models with Right-Censored Event Times. *Biometrical Journal* **48**(6), 1029–1040 (2006)
15. Girdhar, R., El-Nouby, A., Liu, Z., Singh, M., Alwala, K.V., Joulin, A., Misra, I.: ImageBind: One Embedding Space To Bind Them All. pp. 15180–15190 (2023)
16. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative Adversarial Nets. In: *Advances in Neural Information Processing Systems*. vol. 27 (2014)
17. Graf, E., Schmoor, C., Sauerbrei, W., Schumacher, M.: Assessment and comparison of prognostic classification schemes for survival data. *Statistics in Medicine* **18**(17-18), 2529–2545 (1999)
18. Han, W., Chen, H., Poria, S.: Improving Multimodal Fusion with Hierarchical Mutual Information Maximization for Multimodal Sentiment Analysis (Sep 2021), arXiv:2109.00412 [cs]
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)
20. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus) (2023)
21. Huang, Z., Zhan, X., Xiang, S., Johnson, T.S., Helm, B., Yu, C.Y., Zhang, J., Salama, P., Rizkalla, M., Han, Z., Huang, K.: SALMON: Survival Analysis Learning With Multi-Omics Neural Networks on Breast Cancer. *Frontiers in Genetics* **10** (2019)
22. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised Contrastive Learning. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 18661–18673 (2020)
23. Kvamme, H., Borgan, : Continuous and discrete-time survival prediction with neural networks. *Lifetime Data Analysis* **27**(4), 710–736 (Oct 2021)
24. Louis, D.N., Perry, A., Wesseling, P., Brat, D.J., Cree, I.A., Figarella-Branger, D., Hawkins, C., Ng, H.K., Pfister, S.M., Reifenberger, G., Soffietti, R., von Deimling, A., Ellison, D.W.: The 2021 WHO Classification of Tumors of the Central Nervous System: a summary. *Neuro-Oncology* **23**(8), 1231–1251 (Aug 2021)
25. Oord, A.v.d., Li, Y., Vinyals, O.: Representation Learning with Contrastive Predictive Coding (Jan 2019), arXiv:1807.03748 [cs, stat]
26. Sanchez, E.H., Serrurier, M., Ortner, M.: Learning Disentangled Representations via Mutual Information Estimation. In: *Computer Vision – ECCV 2020*. pp. 205–221 (2020)
27. Steyaert, S., Qiu, Y.L., Zheng, Y., Mukherjee, P., Vogel, H., Gevaert, O.: Multimodal deep learning to predict prognosis in adult and pediatric brain tumors. *Communications Medicine* **3**(1), 1–15 (Mar 2023)
28. Tishby, N., Zaslavsky, N.: Deep Learning and the Information Bottleneck Principle (Mar 2015), arXiv:1503.02406 [cs]
29. Vale-Silva, L.A., Rohr, K.: Long-term cancer survival prediction using multimodal deep learning. *Scientific Reports* **11**(1), 13505 (Jun 2021)

30. Wetstein, S.C., de Jong, V.M.T., Stathonikos, N., Opdam, M., Dackus, G.M.H.E., Pluim, J.P.W., van Diest, P.J., Veta, M.: Deep learning-based breast cancer grading and survival analysis on whole-slide histopathology images. *Scientific Reports* **12**(1), 15102 (Sep 2022)
31. Zadeh, A., Chen, M., Poria, S., Cambria, E., Morency, L.P.: Tensor Fusion Network for Multimodal Sentiment Analysis. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pp. 1103–1114. Association for Computational Linguistics (Sep 2017)
32. Zhou, R., Zhou, H., Chen, B.Y., Shen, L., Zhang, Y., He, L.: Attentive Deep Canonical Correlation Analysis for Diagnosing Alzheimer’s Disease Using Multimodal Imaging Genetics. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 681–691. Springer Nature Switzerland (2023)