



OPEN Deep survival analysis in multimodal medical data: a parametric and probabilistic approach with competing risks

Alba Garrido[✉], Alejandro Almodóvar[✉], Patricia A. Apellániz[✉], Santiago Zazo[✉] & Juan Parras[✉]

Accurate survival prediction is critical in oncology for prognosis and treatment planning. Traditional approaches often rely on a single data modality, limiting their ability to capture the complexity of tumor biology. To address this challenge, we introduce a multimodal deep learning framework for survival analysis capable of modeling both single and competing risks scenarios, evaluating the impact of integrating multiple medical data sources on survival predictions. We propose SAMVAE (Survival Analysis Multimodal Variational Autoencoder), a novel deep learning architecture designed for survival prediction that integrates six data modalities: clinical variables, four molecular profiles, and histopathological images. SAMVAE leverages modality-specific encoders to project inputs into a shared latent space, enabling robust survival prediction while preserving modality-specific information. We evaluate SAMVAE on two cancer cohorts—breast cancer and lower-grade glioma—applying tailored preprocessing, dimensionality reduction, and hyperparameter optimization. The results demonstrate the successful integration of multimodal data for both standard survival analysis and competing risks scenarios across different datasets. Our model achieves competitive predictive performance compared to state-of-the-art multimodal survival models. To the best of our knowledge, among publicly available implementations, this framework represents an early approach to integrating clinical, multi-omics, and imaging data for continuous-time competing risks. SAMVAE introduces a probabilistic framework that supports the generation of personalized survival curves from heterogeneous data sources. Its parametric formulation enables the derivation of clinically meaningful statistics from the output distributions, providing patient-specific insights through interactive multimedia. These properties position it as a methodological tool with the potential to support risk modeling and interpretable survival analysis in oncology. Importantly, its main contribution is a unified probabilistic multimodal fusion framework.

Keywords Variational autoencoders, Multimodality, Survival analysis, Competing risks, Personalized medicine

Healthcare data are inherently multimodal, involving diverse sources such as histopathological images, genomic sequencing, and structured clinical records¹. Traditional single-modality approaches, those relying on only one type of data, often fall short of capturing the complexity and richness of medical information². In contrast, multimodal approaches integrate and analyze multiple data types simultaneously, offering a more comprehensive view of patient health and disease mechanisms.

The integration of multimodal data has become essential in oncology research³. Furthermore, studies have increasingly highlighted the critical relationships between histopathological imaging and genetic events⁴. In addition, the adoption of multi-modality approaches is a cornerstone of personalized medicine, as it enables an all-encompassing profile of each patient². This holistic perspective facilitates the development of tailored therapeutic strategies and more precise prognostic assessments, ultimately improving patient outcomes¹.

Despite the promise of multimodal data in the advancement of cancer research, its high dimensionality presents both opportunities and challenges. Molecular and imaging datasets often involve high-dimensional and data-intensive formats; genomic profiles include thousands of features⁵, while Whole Slide Images (WSIs)

Information Processing and Telecommunications Center, ETSI Telecomunicación, Universidad Politécnica de Madrid, Madrid, Spain. ✉email: alba.garrido.lopez@upm.es

used in digital pathology often exceed dimensions of $100,000 \times 100,000$ pixels per slide⁶. These modalities also differ greatly in how they are generated, pre-processed, and represented, making direct integration difficult due to discrepancies in scale, structure, and data format^{1,7}. This curse of dimensionality complicates the analysis and increases the risk of overfitting, particularly when the number of features far exceeds the number of samples.

These challenges are especially critical in the context of medical research, where cancer remains one of the leading causes of death worldwide². A central objective in oncology is to estimate patient survival time, defined as the time until clinically significant events, such as disease progression or death, occur⁸. To address this challenge, survival analysis provides statistical tools for modeling time-to-event data, including scenarios with competing risks, where alternative events may preclude or alter the occurrence of the primary event of interest⁹.

Despite numerous studies that address the problem of survival analysis, many existing approaches face significant limitations. Some rely on oversimplified models that fail to capture the complexity and heterogeneity of cancer data². In contrast, multimodal models present additional limitations such as the lack of interpretability in their results¹⁰, making their integration into clinical decision-making challenging. Furthermore, only a limited number of studies leverage truly multimodal data, and the lack of transparency in the code of some models, coupled with the inaccessibility of certain private datasets, significantly hinders reproducibility and fair benchmarking across the field, as reported by Wiegbe et al.¹⁰.

Building on these motivations, we propose SAMVAE (Survival Analysis Multimodal Variational Autoencoder), a multimodal integration framework for parametric survival analysis and competing risks. To address the inherent heterogeneity across different modalities of data, we adopt a flexible architecture based on Variational Autoencoders (VAEs)¹¹, generative models capable of capturing complex nonlinear relationships within data. This allows us to learn shared latent representations across modalities while enabling interpretability at the output level, as VAEs can model parametric distributions over both covariates and event times, provided the log-likelihood is differentiable¹². This parametric and fully probabilistic formulation is central to our approach, as it enables the principled estimation of uncertainty, allows for flexible distributional assumptions beyond proportional hazards¹³, and provides analytical access to key survival quantities such as event probabilities and cumulative incidence functions. Importantly, our primary objective is to provide a robust multimodal fusion framework, rather than to pursue consistent gains in predictive accuracy alone.

To validate the proposed model, we conduct experiments on breast cancer and lower-grade glioma datasets, leveraging six distinct data modalities. In line with the general classification of oncology data into clinical, molecular, and imaging domains², our study incorporates clinical information, four molecular modalities (DNA methylation (DNAm), Copy Number Variation (CNV), microRNA (miRNA), and RNA sequencing (RNAseq)), and WSIs.

An overview of the proposed architecture is shown in Fig. 1, which outlines the core components of SAMVAE. The model employs modality-specific encoders to transform heterogeneous input data into a common latent space that captures integrated patient representations. This latent space is used for the time prediction, with the specific modeling strategy depending on the nature of the dataset: standard survival analysis is applied when there is a single event of interest, whereas competing risks modeling is used when multiple mutually exclusive events are present. This architecture supports flexible fusion and end-to-end training and optimization, enabling SAMVAE to effectively handle the complexity of multimodal biomedical data across different survival modeling scenarios.

Then, our main contributions consist of:




Fig. 1. General architecture of the proposed SAMVAE framework. The model receives multimodal input data, which are encoded into a shared latent space using modality-specific encoders. This latent representation is used to predict time-to-event outcomes, either survival analysis or competing risks modeling, depending on whether the dataset includes one or multiple mutually exclusive events. Figure created using Canva.

- We introduce SAMVAE, a novel multimodal extension of the SAVAE framework¹², designed to jointly model heterogeneous multimodal biomedical data. By employing modality-specific encoders and projecting them into a shared latent space, SAMVAE enables flexible integration of diverse data types with end-to-end trainability.
- We extend the SAMVAE framework to flexibly support both standard survival analysis and competing risks modeling, enabling task-specific adaptation within a unified multimodal architecture. To the best of our knowledge, among publicly available implementations, this work constitutes one of the few available multimodal approaches capable of jointly modeling clinical, multi-omic, and imaging data for continuous-time competing risks. The proposed framework thus enables personalized survival prediction across heterogeneous data sources (tabular and image data).
- We take advantage of the parametric and probabilistic formulation of SAMVAE to generate individualized, interpretable, and interactive survival curves. This facilitates a deeper understanding of the multimodal factors influencing patient-specific risk trajectories, thereby enhancing the model's clinical transparency and potential for real-world applicability.
- We use publicly available datasets and release our full codebase, promoting transparency, reproducibility, and enabling independent validation.

Background

Survival analysis

Survival Analysis, also known as time-to-event analysis, is a statistical method used to model the time until a specific event occurs. In healthcare, survival analysis plays a crucial role in predicting and understanding health-related events such as disease progression, hospital readmissions, and mortality¹⁴. It supports tasks such as prognostic modeling, in which clinicians assess the likelihood of adverse outcomes for specific patient groups and identify high-risk individuals. A key feature of survival analysis is its ability to handle censored data—instances where the exact time of the event is unknown, either because the event has not yet occurred or the subjects have left a study¹⁵.

We consider that we have a dataset with N observations, each represented by triplets $\mathcal{D} = (x_i, t_i, y_i)_{i=1}^N$. The vector $x_i = (x_i^1, \dots, x_i^L)$ comprises L covariates for individual i , while t_i denotes the time at which an event occurs. The indicator $y_i \in \{0, 1\}$ specifies whether the event was observed ($y_i = 1$) or censored ($y_i = 0$). When an observation is censored, it means that no event was recorded up to time t_i ¹².

The survival function, $S(t | x)$, is a non-increasing function of t that indicates the probability that the event of interest has not yet occurred by time t , conditional on covariates x . It is defined as $S(t | x) = P(T > t | x) = 1 - F(t | x)$ where $F(t | x)$ is the Cumulative Distribution Function (CDF) of event times.

Competing risks

Competing Risks refer to situations in survival analysis where multiple mutually exclusive event types can occur, and the occurrence of one event prevents the occurrence or observation of the others. The Cumulative Incidence Function (CIF) becomes the key quantity of interest in competing risks settings. Let $Y \in \{0, 1, \dots, K\}$ be a discrete random variable, where $Y = 0$ denotes exclusively censoring, and $Y = k$ indicates the occurrence of event type k . The CIF is defined only for the actual event types $k \in \{1, 2, \dots, K\}$, and not for the censored case:

$$\text{CIF}_k(t | x) = P(T \leq t, Y = k | x),$$

where T is the time-to-event, $k \neq 0$, and $Y = k$ indicates the event type. The competing risks model reduces to the standard survival analysis setting when $K = 1$, that is, when $Y \in \{0, 1\}$, meaning there can be either a single event of interest ($Y = 1$) or censoring ($Y = 0$). Unlike the survival function, which aggregates all failure types, the CIF separates cause-specific probabilities. Furthermore, the CIF does not converge to 1 as $t \rightarrow \infty$ and the sum across all risks equals the overall event probability, excluding censoring.

As shown by Apellániz et al.¹⁶, the CIF is defined as:

$$\text{CIF}_k(t | x) = P(T \leq t | Y = k, x) \cdot P(Y = k | x) \quad (1)$$

where the first term (the conditional survival probability for each event type) can be obtained via risk-specific predictors and the second (the risk probability) through a latent-variable classifier. For censored observations, the survival contribution is derived from the fact that no event occurred before the censoring time t , leading to the following definition:

$$S(t | x) = 1 - \sum_{k=1}^K \text{CIF}_k(t | x). \quad (2)$$

Related works

Survival analysis models

Traditional survival analysis has been dominated by the Cox proportional hazards (CoxPH) model¹³, which assumes a linear relationship between covariates and the log-risk function and constant hazard ratios over time. While CoxPH remains popular for its interpretability, its assumptions are often too restrictive for high-dimensional or time-varying data, leading to the development of several deep learning-based extensions. DeepSurv¹⁷ replaces the linear predictor in CoxPH with a deep neural network, enabling the capture of non-

linear covariate interactions while maintaining the proportional hazards assumption. Extensions such as Cox-Time¹⁸ relax this assumption by allowing time-varying effects, while Cox-PASNet¹⁹ integrates prior biological knowledge and regularization for high-dimensional omics data.

Moving beyond the proportional hazards framework, models like DeepHit²⁰ directly estimate the discrete-time survival distribution, allowing for flexible, non-linear, and non-monotonic risk modeling. Its extension, Dynamic-DeepHit²¹, incorporates recurrent neural networks (RNNs) to handle longitudinal data with time-dependent covariates. Fully parametric deep survival models have also gained attention. Approaches such as Deep Weibull²² and Deep Survival Machines (DSM)²³ assume parametric survival distributions and leverage mixture models to increase flexibility. Additionally, SAVAE¹² builds on these parametric advances by introducing a VAE-based framework that models survival times through flexible parametric distributions. Its design enables effective handling of high-dimensional, heterogeneous, and censored data.

Competing risks models

The Fine-Gray (FG) model²⁴ is one of the most established classical approaches for competing risks. By modeling the subdistribution hazard directly, the method enables estimation of the CIF for each event type. However, like CoxPH, FG relies on the proportional hazards assumption and lacks flexibility in complex, real-world scenarios.

Deep learning has introduced more flexible alternatives. DeepHit²⁰ estimates CIFs for multiple events using shared and cause-specific subnetworks, trained via a composite loss. Though powerful, it operates in discrete time and lacks closed-form CIFs. DeepComp²⁵ builds on this by modeling cause-specific discrete hazards with RNNs, improving temporal modeling but still constrained by time discretization. Similarly, CRESA²⁶ uses RNNs to model recurrent CIFs and is particularly suited for sequential data, though it shares the same discrete-time limitations.

To address these, SSMTL²⁷ frames competing risks as a multi-task classification problem, enabling semi-supervised learning but reducing temporal granularity. In contrast, DeepCompete²⁸ introduces a continuous-time framework based on neural ordinary differential equation blocks, offering more precise CIF estimation at the cost of increased computational complexity. CR-SAVAE¹⁶ proposes a generative VAE-based approach that enables accurate cumulative incidence estimation without relying on proportional hazards assumptions, valid for continuous and discrete time. It improves the interpretability of the results by allowing direct analysis of covariate effects and supports robust statistical inference, facilitating personalized and clinically meaningful risk assessment.

State of the art: multimodal survival approaches in medical data

Survival Analysis Models To assess the current state of the field of survival analysis using deep learning, a comprehensive survey¹⁰ published in 2024 initially identified 211 relevant articles in this domain, with an openly maintained codebase to support reproducibility. This review underscores a persistent gap: very few survival models fully exploit multimodal data using public repositories with transparent and standardized evaluation protocols.

During the literature review, several multimodal survival analysis models were identified. Many potentially relevant studies relied on private or restricted datasets (e.g., hospital cohorts), limiting reproducibility and fair benchmarking²⁹, or lacked accessible code despite using public data, as in CapSurv³⁰. Therefore, many were excluded from direct comparison with our approach due to limitations in data types, target populations, or methodological scope.

Our work focuses on integrating clinical data, histopathological images, and multi-omics in specific cancers. The inclusion criteria required that each model incorporate three data modalities (clinical, omic, and images) and use The Cancer Genome Atlas (TCGA) public dataset³¹, one of the most widely adopted and publicly available repositories. From these studies, only 29 addressed multimodality, and a subset of 6 studies utilized the widely recognized and publicly available TCGA repository and made their source code accessible on GitHub.

This subset of six studies, although focused on multimodal integration in oncology, exhibits key limitations that distinguish them from our approach. AdvMIL³² focuses exclusively on histopathological images, without incorporating clinical or omics data. PAGE-Net³³ integrates histological images with gene expression data but is tailored specifically for glioblastoma and relies on pathway-level genomic integration, diverging from our use of raw multi-omics data. GD-Net and ConcatAE/CrossAE^{34,35} incorporate only three omics modalities, omitting both clinical and imaging data, thereby falling short of full multimodal integration. MultiSurv³⁶ supports multiple modalities but is trained across 33 cancer types simultaneously, lacking the tumor-specific optimization that is central to our study. Lastly, BioFusionNet³⁷ employs a highly specialized and computationally intensive framework based on large-scale self-supervised pretraining on over 35 million histopathological patches (DINO and MoCoV3) and dual cross-attention mechanisms, processing up to 500 images per patient. This results in extremely high-dimensional embeddings from clinical, omics, and imaging data, accompanied by substantial computational costs. The results reported for BioFusionNet and its comparative models were obtained by training on a single TCGA dataset, which integrates clinical, omics, and imaging modalities. Therefore, we compare our approach against the reported results of BioFusionNet in the literature and the other models referenced therein.

In addition to BioFusionNet, several recent multimodal survival prediction models have been proposed but are not covered in the survey mentioned¹⁰. MultiDeepCox-SC³⁸ integrates image-derived risk scores with clinical and gene expression data using the Cox proportional hazards model¹³. MCA³⁹ introduces a genomics-guided co-attention mechanism to align whole-slide images with genomic and clinical features. MultiSurv (mentioned above)³⁶ concatenates feature vectors into a unified representation for survival estimation. HFBSurv⁴⁰ leverages a hierarchical attention-based bilinear structure to process data progressively from simple to complex levels. PathomicFusion⁴¹ applies gating-based attention to suppress noise and highlight relevant features. TransSurv⁴² employs cross-attention transformers to combine histopathological, genomic, and clinical inputs, capturing

intermodal dependencies. These approaches highlight the diversity of fusion strategies explored in multimodal survival modeling, with reported results on the TCGA breast cancer cohort available in the BioFusionNet benchmark, which we use as a basis for comparison in our experimental evaluation for survival analysis.

Competing risks models

To evaluate the current state of the field of competing risks using deep learning, the same comprehensive survey mentioned above was also reviewed¹⁰, this time filtering specifically on methods designed to handle competing risks. A few relevant works were identified; however, they present notable limitations. For instance, DeepPAMM⁴³ addresses competing risks and includes a multimodal experiment combining tabular data and point clouds. However, this model operates exclusively in the discrete-time domain because it uses a penalized negative Poisson log-likelihood as a loss function, which requires datasets with a large number of observations, an assumption that does not typically hold in real-world clinical datasets. Neural Fine-Gray⁴⁴ focuses exclusively on tabular datasets and explicitly mentions the extension to other data types, such as time series or images, as future work. Therefore, we exclude it from comparison due to its current lack of support for multimodality. Lastly, SURVNODE⁴⁵ explores multi-state models and competing risks, but its experiments are performed separately on either tabular or time series data, without integrating multiple modalities. These limitations underscore the need for models capable of jointly handling competing risks and multimodal inputs, such as tabular clinical or omic data and imaging, using real-world datasets.

In addition to the models covered in the survey, we identified an additional approach not included in it: Deep-CR MTLR⁴⁶, which includes both clinical and imaging modalities; however, it does not include an omics modality. In addition, this model operates exclusively in the discrete-time domain. This introduces several limitations, most notably, the need for a dataset with a large number of observations, a condition that is often unrealistic in real-world settings. Furthermore, training Deep-CR MTLR requires fitting a separate model for each time point and for each event type, making the approach computationally intensive.

To date, and to the best of our knowledge, within currently available approaches, no previous work has proposed a multimodal deep learning framework that explicitly models competing risks in continuous time. Existing multimodal approaches predominantly rely on conventional Cox proportional hazards models or focus on single-event survival settings, thereby overlooking the complexity introduced by mutually exclusive event types. In contrast, our study presents a publicly validated multimodal framework that integrates competing risks modeling within a unified architecture encompassing clinical, imaging, and multi-omics data, with the flexibility to handle both continuous and discrete time-to-event data. Additionally, our approach is parametric and does not rely on the proportional hazards assumption. This versatility enhances its applicability in real-world healthcare settings, where both data types and event structures are heterogeneous.

Architecture: from unimodal to multimodal VAEs for survival

In this section, we provide a detailed description of the SAVAE model¹² and its extension to competing risks (CR-SAVAE¹⁶), which serve as the foundation for our proposed SAMVAE. While these models were briefly introduced in the related work section, we now detail their structure and survival modeling framework in order to clarify the design choices behind our multimodal extension. Building on this foundation, we then introduce the architecture of SAMVAE, which generalizes the unimodal models to integrate heterogeneous data sources effectively.

Unimodal SAVAE for single and competing risks

SAVAE¹² introduces a generative framework based on VAEs¹¹, allowing it to estimate survival times by parametrically modeling the underlying distribution, whether Weibull, log-normal, or other differentiable forms. This removes the need for strong assumptions, such as proportional hazards, and enables the incorporation of complex, non-linear interactions between covariates and survival outcomes. Moreover, a tailored Evidence Lower BOund (ELBO)¹¹ formulation facilitates the integration of censored data, ensuring robust estimation of time-varying risk factors and covariate dependencies.

CR-SAVAE¹⁶ extends this architecture to the competing risks setting by jointly modeling multiple failure causes through a shared generative framework. Each risk is associated with its own parametric distribution, whose parameters are inferred from a common latent representation modulated by the specific event type.

To clarify the probabilistic structure underlying these models, Fig. 2 presents the neural network architecture used in both SAVAE and CR-SAVAE, adapted from¹⁶. To facilitate understanding of the notation, we provide an appendix summarizing all symbols and definitions used throughout the paper (see Appendix A). The encoder and decoder modules, implemented as multilayer perceptrons (MLPs), are tailored to handle tabular biomedical data. The architecture maps input features into a latent representation, which is then used to reconstruct the input and estimate event times.

The encoder function, denoted by ϕ , maps the input to the latent space. The decoder θ reconstructs the observed covariates from the latent representation. For survival modeling, ψ represents the time predictor(s) that generate the parameters of the time-to-event distribution(s). In the standard survival setting, there is a single-time predictor ψ , as only one type of event is modeled. In contrast, in the competing risks setting, there are multiple time predictors ψ_k —one for each event type—each responsible for modeling the distribution of time to a specific cause of failure. Additionally, ω denotes the event classifier, which predicts the most likely cause of failure among competing risks.

In SAVAE¹², the proposed ELBO formulation is:




Fig. 2. SAVAE and CR-SAVAE architecture. Circles represent latent variables, and rectangles are Neural Networks. Shaded nodes indicate components shared in both SAVAE and CR-SAVAE, while unshaded nodes represent elements specific to the competing risks extension. The parameters ϕ , θ , ψ , and ω denote the weights of the neural networks used. We use a hat ($\hat{\cdot}$) to denote variables predicted by the model.

$$\hat{L}(x, \theta, \psi, \phi) = -\frac{1}{N} \sum_{i=1}^N (-D_{KL}(q_\phi(z|x_i)||p(z)) + \log p_\psi(t_i|g_\phi(x_i, \epsilon_i))) + \sum_{l=1}^L \log p_\theta(x_i^l|g_\phi(x_i, \epsilon_i))) \quad (3)$$

This formulation includes the Kullback-Leibler (KL) divergence term (denoted as D_{KL}). The KL term acts as a regularizer, pushing the approximation of the posterior towards the prior. The latent variable z is sampled using the reparameterization trick as $z = g_\phi(x_i, \epsilon_i)$, which includes both the input data x_i and the noise $\epsilon_i \sim \mathcal{N}(0, I)$, which is a standard Gaussian random variable. The loss also includes the time prediction term, which models the likelihood of the observed time-to-event t_i , and the covariates reconstruction term, which ensures accurate reconstruction of each input feature x_i^l . In the competing risks setting, this loss formulation is extended by maintaining the same KL divergence term and covariate reconstruction loss, but adapting the time-to-event component to account for the presence of multiple possible events. Specifically, CR-SAVAE estimates a separate time distribution for each cause and additionally includes an event prediction term to identify the most likely event type. This event prediction branch was not required in the original SAVAE, where only a single risk was considered.

Following the classical formulation in survival analysis^{47,48}, the contribution of each individual i to the likelihood function under censoring is expressed as:

$$p_\psi(t_i|g_\phi(x_i, \epsilon_i)) = h(t_i|g_\phi(x_i, \epsilon_i))^{y_i} S(t_i|g_\phi(x_i, \epsilon_i)), \quad (4)$$

where y_i is the censoring indicator, defined as $y_i = 0$ if censored and $y_i = 1$ if event experienced. Here, $h(t_i|g_\phi(x_i, \epsilon_i))$ denotes the hazard function, which quantifies the instantaneous rate of event occurrence at time t_i , given that the individual has survived (i.e., has not yet experienced the event) up to that moment.

This formulation ensures that the hazard function term is only considered when the data are not censored ($y \neq 0$), and it has also been employed in more recent frameworks such as SAVAE. In this work, the Weibull distribution is used for modelling time-to-event data, although any time distribution could be used instead, given that its log-likelihood is differentiable.

SAMVAE architecture for single and competing risks

We now introduce SAMVAE, our proposed multimodal extension of the two previously described architectures, designed to address both standard survival analysis (when a single event of interest is present in the dataset) and competing risks scenarios (which involve the presence of multiple mutually exclusive event types). SAMVAE is capable of integrating heterogeneous data modalities, including both tabular data (e.g., clinical and omic modalities) and image data. It leverages a shared latent representation learned from all available modalities, enabling joint modeling of heterogeneous inputs for time-to-event prediction. Figure 3 illustrates the architecture adapted for our proposed SAMVAE framework. Each input modality $x^{(m)}$ is first encoded into a modality-specific latent representation $z^{(m)}$ via a dedicated encoder $\phi^{(m)}$. These latent vectors are subsequently concatenated into a unified latent representation z .

The shaded blue components highlight the novel contributions introduced in SAMVAE, including decoders and predictive branches specifically tailored for temporal and clinical outcomes. Meanwhile, the remaining gray-shaded components correspond to SAVAE and CR-SAVAE preserved from the base architecture.

The proposed model involves as many Deep Neural Networks (DNNs) as the number of encoders and decoders required, depending on the number of modalities involved. An additional module is trained for time-to-event prediction. These decoders are designed to infer covariates, while a dedicated predictor is trained to estimate the time parameters, respectively. Specifically, for M modalities, the total number of DNNs in the model depends on the type of survival task. In a survival analysis scenario, the model uses $2M + 1$ networks: M encoders, M decoders for reconstructing covariates, and one dedicated to time-to-event prediction. In the case




Fig. 3. General architecture of SAMVAE. Gray-shaded nodes denote components shared with both SAVAE and CR-SAVAE. Unshaded nodes are specific to the CR-SAVAE, while the remaining colored nodes represent novel additions introduced by SAMVAE. Each input modality $x^{(m)}$ is encoded via its own encoder of parameters $\phi^{(m)}$ into a latent representation $z^{(m)}$. The concatenated latent vector is used to decode covariates θ , survival time(s) ψ , and (if competing risks) event type ω .

of competing risks, the model requires $2M + K + 1$ networks. This includes M encoders, M covariate decoders, K separate predictors for time-to-event parameters (one predictor per risk), and one additional network acting as a classifier to predict which event (risk) will occur.

In deep learning, architectural choices are guided by data characteristics and the inductive biases of neural networks⁴⁹. Given the multimodal nature of the dataset, comprising tabular (clinical and omic) and image data, distinct encoder-decoder architectures are employed for each modality: MLPs for tabular data and Convolutional Neural Networks (CNNs) for image data.

The latent spaces $z^{(m)}$ generated by each encoder are concatenated into a unified latent representation z , resulting in an approximate posterior $q_\phi(z|x) = q_\phi^{(1)}(z^{(1)}|x^{(1)}) q_\phi^{(2)}(z^{(2)}|x^{(2)}) \dots q_\phi^{(M)}(z^{(M)}|x^{(M)})$, where the joint posterior is defined by the concatenation of all modality-specific latent distributions. This shared latent space z serves as the foundation for generating samples that are fed into all decoders. By ensuring that each decoder receives information from all modalities, the model effectively integrates complementary data sources. This is particularly crucial for the time-to-event predictor(s), as it enables survival analysis and competing risks predictions to be informed by the most relevant modalities, since it estimates the time-to-event distribution parameters from the shared latent representation.

In the present work, the Weibull distribution is adopted for modeling time-to-event data. However, any time distribution with a differentiable log-likelihood could be used instead, including those defined over discrete time. Although the proposed approach is compatible with both continuous and discrete time settings, we restrict our experiments to the continuous Weibull distribution due to its widespread use in survival analysis⁵⁰, its demonstrated effectiveness in prior studies²², and the fact that it does not assume proportional hazards.

Building upon the formulation presented in the Equation 3, where the ELBO explicitly accounts for each dimension of the covariates vector, the proposed extension adapts this framework to handle multimodal data. Specifically, since the model includes as many encoders/decoders as there are modalities, the ELBO is modified to include an additional summation over the number of modalities. This additional summation ensures that each modality contributes appropriately to the overall objective, resulting in the following ELBO formulation (Equation 5). Additionally, SAMVAE operates in a parametric manner: both the time and the covariates can follow any distribution that has a differentiable log-likelihood function, which highlights its flexibility.

$$\hat{L}(x, \theta, \psi, \phi) = -\frac{1}{N} \sum_{i=1}^N \left(-\text{D}_{\text{KL}}(q_\phi(z|x_i)||p(z)) + \log p_\psi(t_i|g_\phi(x_i, \epsilon_i)) \right. \\ \left. + \sum_{m=1}^M \sum_{l_m=1}^{L_m} \log p_\theta(x_i^{l_m}|g_\phi(x_i, \epsilon_i)) \right) \quad (5)$$

In this formulation:

- $\text{D}_{\text{KL}}(q_\phi(z|x_i)||p(z))$: This term represents the KL divergence between the approximate posterior $q_\phi(z|x_i)$ and the prior distribution $p(z)$, where z denotes the fused latent representation obtained from the concatenation of all modality-specific latent vectors $z^{(m)}$. As a regularization term, it encourages the aggregated latent variable z to remain close to the prior distribution.
- $\log p_\psi(t_i|g_\phi(x_i, \epsilon_i))$: This term, also known as the time prediction term, corresponds to the likelihood of the time-to-event variable t_i , given the deterministic transformation $g_\phi(x_i, \epsilon_i)$, which includes both the input

data x_i and the noise ϵ_i . The function g_ϕ is defined as the concatenation of all the $g_\phi^{(M)}$, one for each modality M .

- $\sum_{m=1}^M \sum_{l_m=1}^{L_m} \log p_\theta(x_i^{l_m} | g_\phi(x_i, \epsilon_i))$: This summation accounts for the reconstruction term of the covariates vector $x_i^{l_m}$, ensuring that each dimension l_m is explicitly modelled. Here, M represents the number of modalities, and L_m denotes the dimensions of the covariates specific to each modality m .

Similarly, the total loss used for training SAMVAE is adapted to account for multiple event types. In this setting, the censoring indicator is defined as $y_i = 0$ for censored patients, and $y_i \neq 0$ for uncensored patients who experienced event type k . This generalization allows the model to distinguish between competing outcomes during training and adjust the loss function accordingly. Therefore, in SAMVAE, we define the loss as:

$$\begin{aligned} \hat{\mathcal{L}}(x, \theta, \psi, \phi, \omega) = & -\frac{1}{N} \sum_{i=1}^N \left[-D_{\text{KL}}(q_\phi(z | x_i) \| p(z)) + \sum_{m=1}^M \sum_{l_m=1}^{L_m} \log p_\theta(x_i^{l_m} | g_\phi(x_i, \epsilon_i)) \right. \\ & + \mathbb{I}(y_i \neq 0) \left(\log p_\psi(t_i | g_\phi(x_i, \epsilon_i)) + \log p_\omega(y_i | g_\phi(x_i, \epsilon_i)) \right) \\ & \left. + \mathbb{I}(y_i = 0) \log \left(1 - \sum_{k=1}^K \text{CIF}_k(t_i | g_\phi(x_i, \epsilon_i)) \right) \right] \end{aligned} \quad (6)$$

This formulation extends the original CR-SAVAE loss by incorporating terms specific to time-to-event modeling and competing risks, while accounting for the joint contribution of multiple data modalities in the latent representation. Once the conditional survival functions $P(T \leq t | Y = k, x)$ and risk probabilities $P(Y = k | x)$ are estimated, the CIFs can be computed as in Eq. 1.

To ensure that the latent representation z is meaningful, we include a reconstruction term in the loss function. While it would be possible to rely solely on the encoder to obtain z for downstream survival prediction, the inclusion of a decoder and the associated reconstruction loss imposes an additional constraint that enforces z to serve as a compact and informative summary of each patient's multimodal data.

In this formulation:

- **KL Divergence Term** $D_{\text{KL}}(q_\phi(z | x_i) \| p(z))$ penalizes deviations between the approximate posterior and the prior over the latent variables, encouraging a regularized and informative latent space.
- **Reconstruction Term** $\sum_{m=1}^M \sum_{l_m=1}^{L_m} \log p_\theta(x_i^{l_m} | g_\phi(x_i, \epsilon_i))$ captures the multimodal reconstruction loss, where each modality m and its L_m constituent features are reconstructed from the shared latent representation z_i . This term ensures that the model learns meaningful representations of all input modalities.
- **Event-Specific Likelihood (Uncensored Cases)** For patients who experienced an event ($y_i \neq 0$), the loss includes two terms: $\log p_\psi(t_i | g_\phi(x_i, \epsilon_i))$ models the time-to-event given that the event type information provided by the latent representation, and $\log p_\omega(y_i | g_\phi(x_i, \epsilon_i))$ estimates the probability of observing the specific event type. These terms guide the model to predict both when and what kind of event is likely to occur.
- **Censoring Term (Censored Cases)** For censored patients ($y_i = 0$), the loss includes the logarithm of the survival probability across all event types: $\log \left(1 - \sum_{k=1}^K \text{CIF}_k(t_i | g_\phi(x_i, \epsilon_i)) \right)$. This term ensures that the model correctly accounts for censoring by maximizing the likelihood of survival beyond the observed time.

For the reconstruction of tabular data, the same log-likelihood functions as those used in SAVAE are employed, corresponding to Gaussian, Bernoulli, and Categorical distributions depending on the variable type. In the case of image data, we use the Mean Squared Error (MSE), which corresponds to a Gaussian likelihood with fixed variance. The total reconstruction loss is obtained by summing the losses across modalities. This additive formulation could be further improved by introducing modality-specific weights, allowing the model to balance their relative contributions more effectively, as suggested in⁵¹.

Experimental analysis

This section introduces the experimental validation framework designed to assess the performance of the proposed SAMVAE model. The goal is to determine its ability to integrate heterogeneous medical data for survival prediction—both in standard survival analysis and competing risks—and to compare it with baseline and state-of-the-art methods under fair conditions.

We work with multimodal patient data that includes tabular information (clinical variables and omics data) and histopathological images in the form of patch-level representations. Each modality differs in dimensionality and data structure.

The ultimate goal of this study is to compare the proposed SAMVAE model with baseline and state-of-the-art approaches, including configurations based on clinical, omics, and imaging modalities. Before performing such comparisons, we first identify the optimal configuration of SAMVAE for each cancer type and data modality. Since the encoders produce a Gaussian latent space, the dimensionality of this latent representation and the number of neurons in each hidden layer constitute the key hyperparameters to be tuned. We perform a grid search over these hyperparameters by exploring multiple configurations, aiming to identify the setup that maximizes model performance.

After selecting the optimal parameters for each modality, we evaluate different multimodal combinations—including various omics modalities and varying numbers of histopathological patches per patient—to determine the best integration strategy. The final selected configuration, integrating clinical, omics, and imaging data, is then used in the main experiments and comparative evaluation against baseline and state-of-the-art methods.

Datasets and preprocessing

TCGA provides a comprehensive repository of genomic and clinical data from over 11,000 patients across 33 cancer types³¹. In this study, we selected two representative cohorts, Breast Invasive Carcinoma (BRCA) and Lower Grade Glioma (LGG), due to the availability of multimodal and publicly accessible patient data. For each cohort, we used two datasets: one for standard survival analysis with a single risk (BRCA-SA and LGG-SA) and another for competing risks modeling (BRCA-CR and LGG-CR).

For the competing risks analysis, we utilized the time-to-event and event status annotations provided for the Progression-Free Interval (PFI) endpoint, as recommended by TCGA guidelines⁵². PFI captures the duration from the date of diagnosis until the first occurrence of a new tumor event, which includes progression, recurrence, or death with evidence of disease. Specifically, it distinguishes between three types of outcomes: 1 for the event of interest (tumor progression or cancer-related death), 2 for competing risks (death without prior progression), and 0 for censored observations. In this study, we applied this competing risks framework to the BRCA-SA and LGG-SA cohorts, utilizing their respective competing risks datasets (BRCA-CR and LGG-CR).

A total of six data modalities were used: (i) clinical variables (tabular), (ii) DNAm, (iii) CNV, (iv) miRNA, (v) RNAseq, and (vi) WSIs. These modalities span three major categories of healthcare data—clinical, molecular, and imaging—commonly employed in multimodal biomedical studies².

Tailored preprocessing strategies were applied to each data modality, taking into account its specific characteristics and underlying data structure. For omics data, due to their high dimensionality, reduction techniques such as the minimal-redundancy-maximal-relevance (mRMR) criterion⁵¹, Principal Component Analysis (PCA), retaining the top 100 components, and filtering by the most frequently mutated genes were applied, following practices established in works like ConcatAE/CrossAE⁵³ and Multisurv⁵⁴. For WSIs, we used the Clustering-constrained Attention Multiple Instance Learning (CLAM) framework⁵⁵, which provides predefined preprocessing configurations for TCGA data for the tissue segmentation and extraction of the most informative patches. Each WSI is divided into smaller patches, which are subsequently ranked by relevance using CLAM's attention-based mechanism. For each patient, we select the top informative patches and use them as input to our model.

Given the inherent heterogeneity of these modalities, each modality required specific preprocessing steps (fully documented in the GitHub repository⁵⁶). However, our primary focus lies in designing a flexible and scalable multimodal integration strategy rather than in preprocessing itself. Although the current preprocessing pipeline is effective, it remains adaptable to potential future improvements.

Performance metrics

In survival analysis, the Concordance Index (C-index) is a commonly used metric to measure the quality of a predictive model in ranking individuals by their risk of experiencing an event (e.g., death, disease recurrence). Formally, for a given time t and two patients i, j , the time-dependent C-index can be expressed as⁵⁷:

$$C_{\text{index}}(t) \approx \frac{\sum_{i \neq j} A_{i,j} \cdot \mathbb{I}(F(t | x_i) > F(t | x_j))}{\sum_{i \neq j} A_{i,j}(t)} \quad (7)$$

where $F(t | x)$ is the estimated cumulative distribution function for time-to-event and $A_{i,j}(t) = \mathbb{I}(Y_i = 1 | T_i < T_j, T_i \leq t)$. A C-index close to 1 indicates better alignment between predicted risk and observed outcomes, while a value of 0.5 corresponds to random predictions.

In competing risks settings, the C-index for event k measures the concordance between the predicted CIF for a patient at the time of the event and that of patients who have not yet experienced the event, considering event timing and censoring. Following the definition in¹⁶, it is computed as:

$$C_{\text{index}}(t) \approx \frac{\sum_{i \neq j} A_{k,i,j} \cdot \mathbb{I}(\text{CIF}_k(t | x_i) \leq \text{CIF}_k(t | x_j))}{\sum_{i \neq j} A_{k,i,j}}, \quad (8)$$

where $A_{k,i,j} = \mathbb{I}(Y_i = K | T_i < T_j, T_i \leq t)$ is an indicator function for whether the pair of patients (i, j) is comparable for event k , and $\mathbb{I}(\cdot)$ is the indicator function.

The IBS quantifies the squared error between predicted survival probabilities and actual outcomes, accounting for censoring via Inverse Probability of Censoring Weighting (IPCW)⁵⁸. The Brier Score (BS) at time t is defined as in¹²:

$$BS(t) = \frac{1}{N} \sum_{i=1}^N \left[\frac{(S(t | x_i))^2}{G(t_i)} \cdot \mathbb{I}(t_i < t, y_i = 1) + \frac{(1 - S(t | x_i))^2}{G(t)} \cdot \mathbb{I}(t_i \geq t) \right], \quad (9)$$

where $S(t | x_i)$ is the predicted survival probability, $\mathbb{I}(\cdot)$ is an indicator function, and $G(t)$ is the censoring survival function. $G(t)$ represents the estimated probability of not being censored at time t , and is obtained using the Kaplan-Meier (KM) estimator⁴⁷.

The BS reflects both discrimination and calibration. To summarize performance over time, the Integrated Brier Score (IBS) averages the BS across the follow-up period as:

$$IBS(t_{\max}) = \frac{1}{t_{\max}} \int_0^{t_{\max}} BS(t) dt, \quad (10)$$

with $t_{\max}$ denoting the maximum follow-up time.

As discussed in¹⁶, in the context of competing risks, the IBS quantifies the overall discrepancy between the predicted and observed CIFs over time, while accounting for the presence of multiple events. The time-dependent Brier Score $BS(t)$ is given by:

$$BS_k(t) = \frac{1}{N} \sum_{i=1}^N \left[\frac{(1 - CIF_k(t | x_i))^2}{G(t_i)} \cdot \mathbb{I}(t_i < t, y_i \neq 0) + \frac{(CIF_k(t | x_i))^2}{G(t)} \cdot \mathbb{I}(t_i \geq t) \right]. \quad (11)$$

Experimental setup for parameter and modality selection

For each modality, the optimal parameter combination was identified using five random seeds to ensure the robustness and stability of the selection process. The optimized parameters were the latent dimension and the hidden size (values tested for each modality and configuration are detailed in the Appendix B, see Table 4). Once the best set of parameters was found, we evaluated the combination of omics and clinical data that yielded the best performance. Similarly, results were analyzed using the combination of clinical data with a different number (1, 5, 10, or 15) of the most informative patches from WSIs per patient, extracted using the CLAM model.

For the final stage of our experimental pipeline—after identifying the best omics configuration, optimal patch selection, and incorporating clinical data—we conducted training using 10 different random seeds to ensure the robustness and reliability of the results. These are the experiments presented in this paper as the main evaluation of the proposed SAMVAE model.

Additionally, we used K-fold cross-validation to select the two hyperparameters: latent dimension and hidden size. Due to the limited number of patients in the datasets, all available data were utilized for both the hyperparameter optimization and the final experiments. During the cross-validation process, we separated the data into 5 folds, where each fold was used once as a validation set while the others were used for training. Within each cross-validation iteration, the validation fold was used for model checkpoint selection and performance evaluation, and hyperparameter configurations were selected based on average performance across all folds.

Given the sensitivity of VAE architectures to initialization conditions, our model includes specific design choices to ensure robust and representative evaluation. Following the methodology discussed in⁵⁹, for models sensitive to initialization, the C-index and IBS are averaged over the three best-performing seeds, consistent with the experimental methodology reported in¹². To ensure full transparency and to provide a complete view of the variability across runs, the results for all 10 random initializations are reported in supplementary material available in the repository⁵⁶. Although the number of seeds is limited, statistical comparison is performed on averaged performance metrics across independent random initializations.

These C-index and IBS values are then used to calculate p -values by performing a unilateral t-test, ensuring statistical significance in the results. Under the common assumption of approximate normality of performance estimates, the t-test provides a standard and widely adopted approach for comparing mean performance values. This choice is motivated by the directional nature of our hypothesis, as we aim to assess whether a given model improves performance over a reference model rather than detecting arbitrary differences. For the C-index, the alternative hypothesis assumes higher values, while for the IBS it assumes lower values. Under this framework, one-sided tests provide a more appropriate and statistically powerful assessment.

The final selection of the best hyperparameter combination is based on a two-stage criterion:

- We applied the Holm-Bonferroni (HB)⁶⁰ correction to control for Type I errors due to multiple testing. This method is more powerful than the traditional Bonferroni correction and is well-suited to our analysis as it does not assume independence between tests. Furthermore, it is widely recommended in medical research for controlling the family-wise error rate in multivariable settings while preserving statistical efficiency⁶¹. Combinations in which the corrected p -values of the C-index and IBS do not exceed the threshold of 0.01 are discarded, thereby applying a conservative selection criterion aimed at reducing potential Type I error inflation due to multiple comparisons, and retaining only statistically robust models.
- Among the remaining combinations, the one that maximizes the CI-IBS (C-index minus Integrated Brier Score) metric is chosen. This metric is selected because a higher C-index indicates better performance, while a lower IBS is preferable; therefore, the IBS is subtracted. By combining them into a single metric, we obtain a balanced assessment of the model's performance.
- In the case of competing risks datasets, the combined metric is computed as the average of the C-index values across all considered risks, minus the average of their respective IBS values. This metric is used to identify the hyperparameter combination that optimizes overall performance across risks.

Therefore, once the best hyperparameters have been identified for each dataset and scenario, and the optimal combination of omics data and image patch number has been determined, these are combined with the clinical modality to establish the final configuration used for the evaluation of SAMVAE and its comparison with state-of-the-art methods.

Results

The experiments reported in this section correspond to the configurations that combine the inclusion of clinical data, the best-performing omics modality (or modalities), and the optimal number of image patches. These configurations represent the most effective fusion strategies identified through hyperparameter tuning and intermediate experiments conducted for each dataset. These final multimodal combinations form the basis for evaluating the performance of SAMVAE.

For the experiments, we employed K-fold cross-validation with 5 folds, using all available patient data and 10 random seeds for consistency. In line with our model’s sensitivity to initialization conditions, we averaged the C-index and IBS over the three best-performing seeds, calculating p-values via a unilateral t-test to ensure statistical significance.

Multimodal survival analysis

This section presents the results obtained from the survival analysis tasks in the LGG-SA and BRCA-SA datasets. In addition, the performance of SAMVAE on BRCA-SA is compared with state-of-the-art approaches presented in Section 2.3.3 for multimodal integration of clinical, omics, and imaging data. Table 1a presents the results of the final multi-modal experiments for the LGG-SA cohort. It compares different combinations of clinical, DNAm, and histological image data (using 10 patches) in terms of survival prediction performance. The selection of DNAm, as the omic modality, and the use of 10 patches for histological data is based on intermediate experiments (detailed in the previous section and exemplified by the tables shown in Appendix C).

While clinical data alone achieves strong predictive performance with a C-index of 0.773 ± 0.089 , the combination with 10 patches preserves this discriminative power (C-index = 0.773 ± 0.043) and simultaneously improves calibration, as indicated by a lower IBS. Moreover, adding both omics and image modalities results in a slightly lower IBS (0.186 ± 0.023), showing consistent gains in model calibration through multimodal integration. Notably, the combinations using all three modalities, clinical + patches, clinical + DNAm, and clinical data alone, exhibit similar overall performance, suggesting that multiple configurations can achieve comparable predictive accuracy and calibration.

Similarly to the experiments conducted for the LGG-SA cohort, the final multimodal configurations for the BRCA-SA dataset are presented in Table 1b. The table evaluates different combinations of clinical, CNV, and histological image data (using 15 patches) in terms of survival prediction performance (choice supported by intermediate experiments, see Appendix C). As with LGG-SA, clinical variables of BRCA-SA alone achieve strong predictive performance (C-index = 0.722 ± 0.057). It is crucial to highlight that both unimodal clinical models and their combinations with other modalities consistently yield strong predictive performance. This observation underlines the value of clinical variables as a robust baseline in survival analysis.

Established clinical evidence identifies traditional prognostic factors—particularly age—as key determinants of survival outcomes⁶², a pattern that is reaffirmed in our results, where clinical variables clearly dominate survival prediction. Therefore, the proposed multimodal integration can be considered successful, as it not only

Clinical	DNA Methylation	10 Patches	C-index	IBS	CI - IBS	p-CI HB	p-IBS HB
(a) Performance metrics for different modality combinations on LGG-SA							
•			0.773 ± 0.089	0.201 ± 0.035	0.572 ± 0.084	1.000	0.591
	•		0.615 ± 0.042	0.196 ± 0.014	0.420 ± 0.049	< 1e ^{−3}	0.639
		•	0.592 ± 0.029	0.195 ± 0.015	0.397 ± 0.037	< 1e ^{−3}	0.773
•	•		0.754 ± 0.050	0.190 ± 0.027	0.565 ± 0.060	1.000	1.000
•		•	0.773 ± 0.043	0.187 ± 0.020	0.586 ± 0.050	1.000	1.000
	•	•	0.620 ± 0.059	0.198 ± 0.019	0.421 ± 0.071	< 1e ^{−3}	0.447
•	•	•	0.761 ± 0.047	0.186 ± 0.023	0.576 ± 0.052	1.000	1.000
Clinical	CNV	15 Patches	C-index	IBS	CI - IBS	p-CI HB	p-IBS HB
(b) Performance metrics for different modality combinations on BRCA-SA							
•			0.722 ± 0.057	0.175 ± 0.025	0.546 ± 0.066	1.000	1.000
	•		0.636 ± 0.034	0.195 ± 0.020	0.441 ± 0.031	< 1e ^{−3}	0.099
		•	0.615 ± 0.074	0.203 ± 0.029	0.411 ± 0.085	0.001	0.034
•	•		0.705 ± 0.061	0.191 ± 0.035	0.514 ± 0.077	1.000	0.602
•		•	0.699 ± 0.046	0.192 ± 0.019	0.507 ± 0.054	0.922	0.195
	•	•	0.640 ± 0.060	0.196 ± 0.021	0.444 ± 0.059	0.003	0.094
•	•	•	0.688 ± 0.060	0.192 ± 0.019	0.496 ± 0.057	0.504	0.178


Table 1. Both tables report the C-index, IBS, and their difference (CI-IBS), expressed as mean ± standard deviation aggregated over the three best-performing seeds within each cross-validation fold. P-values, corrected using HB, are computed by comparing each combination against the best-performing one in terms of C-index and IBS, respectively. Bold values indicate combinations whose performance is not significantly different (*p*-values > 0.01) from the best one, for both C-index and IBS. Detailed results for all seeds and 95% confidence intervals are available in the supplementary material⁵⁶.

reaches the strong discriminative performance of clinical data alone but also improves calibration metrics in some combinations.


To illustrate the performance of our SAMVAE model in the context of survival analysis in LGG-SA and BRCA-SA cohorts, Fig. 4a and b show the average predicted survival curves obtained from different combinations of data modalities, as indicated in the legend. The shaded gray area corresponds to the KM estimate⁴⁷, which serves as a non-parametric reference survival curve calculated from the observed event data. For each modality combination, the survival curve is computed as the mean of the predicted survival probabilities across all test patients. The shaded region around each curve represents the 95% confidence interval.

As shown in Fig. 4a and b, all modality combinations—including those that exclude clinical data—yield survival curves that closely follow the KM estimates for both LGG-SA and BRCA-SA cohorts. This consistent alignment across diverse input configurations (unimodal and multimodal combinations) underscores the model's robustness and its ability to generate accurate and well-calibrated survival predictions, regardless of the modality combination. The reliable overlap between predicted and KM survival curves demonstrates that the proposed framework can effectively and successfully capture the fundamental structure of survival events from heterogeneous and multiple modalities.

The project repository includes an interactive HTML visualization that compares the survival curves predicted by different trained models against the Kaplan-Meier estimate for LGG and BRCA, allowing users to assess model performance across patient groups visually (see Appendix D.1, available at [Interactive LGG Survival Plot](#) and [Interactive BRCA Survival Plot](#)).



(a) Predicted survival curves on the LGG-SA dataset.



(b) Predicted survival curves on the BRCA-SA dataset.

Fig. 4. Average predicted survival curves derived from different combinations of input modalities. Each colored line corresponds to a specific modality configuration, as indicated in the legend, with shaded bands indicating 95% confidence intervals. The shaded gray region represents the KM estimate with a 95% confidence interval. The consistent overlap observed across all modality combinations highlights the model's effectiveness in generating accurate survival predictions.

Comparison with state-of-the-art multimodal models

Some state-of-the-art approaches, such as⁵⁴, report only multimodal results that include clinical data, omitting scenarios in which only omic or imaging data are available, or in which clinical information is absent. This lack of unimodal and ablation baselines limits the interpretability of the reported performance and complicates a precise understanding of the individual contribution of each modality. Our work addresses this gap by providing a transparent and systematic evaluation across all modality combinations, demonstrating that clinical data alone already achieves strong performance, a necessary reference point to assess the added value of multimodal integration fairly.

Table 2 presents the performance comparison of several multimodal models for BRCA-SA prediction using clinical, omics, and histopathology image data. The results for all models, except for SAMVAE, are taken directly as reported in the BioFusionNet study. To assess statistical significance, we performed pairwise independent two-sample t-tests based on summary statistics, implemented using the SciPy library⁶³ and applied the HB correction. Each test evaluates whether there is a statistically significant difference between a given model and the best-performing one.

While BioFusionNet reports the highest mean C-index, the difference compared to our SAMVAE model is not statistically significant ($p = 0.705$), highlighting the competitiveness of our approach. Unlike BioFusionNet, which relies on multiple pretrained models and incurs high computational costs, SAMVAE is trained end-to-end without the need for multiple pretrained components, thereby reducing dependency on the alignment of data modalities with pretrained features. Therefore, we can conclude that SAMVAE is competitive with the state of the art, showing no statistically significant difference with the best-performing model in terms of C-index.

Multimodal competing risks

This section presents the results obtained from competing risks cohorts (LGG-CR and BRCA-CR datasets). As detailed in Section 2.3.3, no multimodal models combining clinical, omics, and imaging modalities, while also modeling continuous time, were found for competing risks. Therefore, we do not compare the results obtained with SAMVAE to the state of the art, and instead present only the ablation study results on the integration of tabular and imaging modalities.

The results of the competing risks analysis for the LGG-CR cohort are presented in Table 3a, showcasing the performance of our multimodal model across different data modalities. In this case, the omics combination that achieved the best results in intermediate experiments (Appendix C) consisted of DNAm, CNV, and miRNA data. Regarding imaging, the best performance was obtained when using five patches per patient. These selections were made by optimizing the average CI-IBS score across the two competing risks, ensuring robust performance across both risks simultaneously rather than treating them independently. The CI-IBS metric is defined as the average C-index minus the average IBS across the two risks. In the LGG-CR setting (see Table 3a), most combinations of modality appear in bold, indicating statistically significant results (p -value > 0.01). Additionally, the IBS tends to be higher for Risk 2, likely due to the smaller number of patients who experience this event compared to Risk 1.

In BRCA-CR (see Table 3b), the clinical modality alone achieved the highest performance. Regarding omic data, the best-performing combination was CNV + RNA-Seq. By adding histological data (10 patches), the C-index obtained for Risk 1 matched that of the clinical data alone, while achieving a significantly lower IBS. All combinations are highlighted in bolditalics, and similarly to the LGG-CR cohort, we observe that the IBS is higher for Risk 2, which can be explained by the lower number of patients experiencing this event compared to Risk 1.

To better visualize this, we illustrate the CIFs obtained with each combination of modalities for each competing risk and the Aalen-Johansen (AJ) estimator⁶⁴, for the BRCA-CR cohort (Fig. 5a and b) and the LGG-CR cohort (Fig. 6a and b).

Unlike the KM estimator, which treats competing events as censored and thus tends to overestimate the probability of each individual event, the AJ estimator properly accounts for the presence of competing risks⁶⁵, providing a more realistic and accurate estimate of the event-specific survival probability. Figure 5a and b show that the predicted CIFs produced by our models for BRCA-competing risks exhibit strong overlap with the AJ estimates, indicating that, at the population level, the models accurately capture the dynamics of cumulative

Model	C-index (Mean ± Std)	p-value HB
SAMVAE	0.699 ± 0.065	0.705
MultiSurv ⁵⁴	0.626 ± 0.072	0.057
MultiDeepCox-SC ³⁸	0.604 ± 0.088	0.072
HFBSurv ⁴⁰	0.536 ± 0.073	0.004
PathomicFusion ⁴¹	0.516 ± 0.080	0.005
MCAT ³⁹	0.700 ± 0.043	0.366
TransSurv ⁴²	0.686 ± 0.047	0.207
BioFusionNet ³⁷	0.768 ± 0.051	1.000

Table 2. Comparison of multimodal models for BRCA-SA prediction using the C-index. Pairwise t-tests were performed to assess significance against the best-performing model, BioFusionNet. We correct p-values using the HB method.

Clinical	DNAm + CNV + miRNA	5 Patches	Risk	C-index	IBS	CI - IBS	p-CI HB	p-IBS HB
(a) Performance metrics for different modality combinations on the LGG-CR								
•			1	0.610 ± 0.044	0.240 ± 0.046	0.384	1.000	0.221
•			2	0.616 ± 0.045	0.217 ± 0.018	0.384	1.000	1.000
	•		1	0.585 ± 0.028	0.230 ± 0.047	0.352	0.593	0.867
	•		2	0.590 ± 0.032	0.240 ± 0.039	0.352	0.621	0.386
		•	1	0.585 ± 0.037	0.227 ± 0.042	0.357	0.807	0.952
		•	2	0.595 ± 0.064	0.239 ± 0.042	0.357	1.000	0.645
•	•		1	0.580 ± 0.024	0.221 ± 0.039	0.344	0.266	1.000
•	•		2	0.589 ± 0.031	0.260 ± 0.045	0.344	0.528	0.025
•		•	1	0.596 ± 0.032	0.204 ± 0.039	0.369	1.000	1.000
•		•	2	0.587 ± 0.020	0.240 ± 0.021	0.369	0.285	0.028
	•	•	1	0.579 ± 0.024	0.232 ± 0.057	0.341	0.211	0.924
	•	•	2	0.591 ± 0.035	0.254 ± 0.022	0.341	0.737	< 1e⁻³
•	•	•	1	0.577 ± 0.022	0.233 ± 0.040	0.323	0.146	0.449
•	•	•	2	0.570 ± 0.027	0.268 ± 0.047	0.323	0.021	0.009
Clinical	CNV + RNAseq	10 Patches	Risk	C-index	IBS	CI - IBS	p-CI HB	p-IBS HB
(b) Performance metrics for different modality combinations on the BRCA-CR								
•			1	0.664 ± 0.040	0.199 ± 0.035	0.471	1.000	0.806
•			2	0.682 ± 0.042	0.205 ± 0.033	0.471	1.000	1.000
	•		1	0.618 ± 0.055	0.195 ± 0.028	0.433	0.296	1.000
	•		2	0.632 ± 0.053	0.190 ± 0.031	0.433	0.074	1.000
		•	1	0.638 ± 0.037	0.206 ± 0.030	0.433	1.000	0.190
		•	2	0.635 ± 0.039	0.200 ± 0.022	0.433	0.035	1.000
•	•		1	0.659 ± 0.055	0.177 ± 0.036	0.456	1.000	1.000
•	•		2	0.632 ± 0.060	0.200 ± 0.035	0.456	0.112	1.000
•		•	1	0.657 ± 0.060	0.196 ± 0.023	0.446	1.000	0.694
•		•	2	0.639 ± 0.026	0.207 ± 0.021	0.446	0.024	0.667
	•	•	1	0.641 ± 0.038	0.204 ± 0.033	0.433	1.000	0.357
	•	•	2	0.639 ± 0.044	0.210 ± 0.042	0.433	0.088	1.000
•	•	•	1	0.666 ± 0.066	0.183 ± 0.037	0.458	1.000	1.000
•	•	•	2	0.637 ± 0.031	0.205 ± 0.032	0.458	0.024	1.000


Table 3. Both tables report the C-index, IBS, and their difference (CI-IBS), expressed as mean ± standard deviation aggregated over the three best-performing seeds within each cross-validation fold. P-values, corrected using HB, are computed by comparing each combination against the best-performing one in terms of C-index and IBS (for each risk), respectively. Bold values indicate combinations whose performance is not significantly different ($p > 0.01$) from the best one. Detailed results for all seeds and 95% confidence intervals are available in the supplementary material⁵⁶.

incidence. The only noticeable deviation appears for Risk 2 in the omics-only configuration, which slightly overestimates the incidence. In contrast, the rest of the modality combinations align closely with the AJ curve. Additionally, we observe that the confidence intervals around the AJ estimator widen considerably after year 10, reflecting increasing uncertainty due to the reduced number of observed events over time. This highlights the limitations of empirical estimators in the presence of data sparsity and reinforces the relevance of model-based approaches under such conditions.


Figure 6a and b show that, for the LGG-CR cohort, the predicted CIFs align closely with the AJ estimates across all modality configurations. For Risk 1, the curves exhibit a strong and consistent overlap with the AJ estimator. In the case of Risk 2, however, we observe considerably wider confidence intervals around the AJ curve, particularly after year 5. This is a consequence of the substantially smaller number of events associated with Risk 2 compared to Risk 1, which increases statistical uncertainty.

Overall, these findings demonstrate a successful integration of multimodal data for competing risks modeling. The fact that multimodal combinations do not underperform compared to clinical data alone, as consistently observed in Tables 3a and 3b, reinforces the complementary value of omic and histological information.

To the best of our knowledge, among publicly available implementations, there are currently no existing multimodal deep learning frameworks that address competing risks modeling in continuous time (rather than



(a) Average predicted CIFs for Risk 1 in the BRCA-CR dataset.



(b) Average predicted CIFs for Risk 2 in the BRCA-CR dataset.


Fig. 5. Average predicted CIFs for the two competing risks in the BRCA-CR dataset. Each colored line corresponds to a different combination of input modalities used by the model, with shaded bands indicating 95% confidence intervals. The predictions are compared against the AJ estimator (shaded region represents a 95% confidence interval), which properly handles competing risks. The close overlap with the AJ curves indicates accurate population-level estimation.

just discrete time) to specific events, while simultaneously integrating clinical, omics, and imaging modalities, as detailed in Section 2.3.3. As a result, direct comparisons with state-of-the-art models are not feasible.


Personalized medicine

One of the key advantages of the proposed approach is its ability to generate personalized survival curves for each patient, accounting for multiple plausible scenarios by sampling from the latent space. This allows the model to capture individual uncertainty and generate distributions over survival times, offering a richer and more informative prognosis than point estimates. These individualized trajectories could provide complementary insights in clinical research settings, exploring factors such as treatment intensity, follow-up frequency, or eligibility for clinical trials.

To illustrate the personalized nature and practical utility of our approach, Fig. 7 presents a comparison of the survival curves predicted by our model for the patients with the best and worst prognoses in the validation set for LGG, alongside the empirical KM estimator for the entire cohort. For each selected patient, we generated 100 survival trajectories by sampling from the posterior distribution of the latent space. The mean survival curve is shown with a 90% confidence band (5th–95th percentiles). The patient with the poorest prognosis shows a steep decline in survival probability shortly after diagnosis, while the best prognosis patient maintains a high probability of survival throughout the follow-up period.



(a) Average predicted CIFs for Risk 1 in the LGG-CR dataset.



(b) Average predicted CIFs for Risk 2 in the LGG-CR dataset.

Fig. 6. Average predicted CIFs for the two competing risks in the LGG-CR dataset. Each colored line corresponds to a different combination of input modalities used by the model, with shaded bands indicating 95% confidence intervals. The predictions are compared against the AJ estimator (shaded region represents a 95% confidence interval), which properly handles competing risks. The close overlap with the AJ curves indicates accurate population-level estimation.

The project repository includes an interactive HTML visualization that allows users to explore the clinical features associated with each patient's survival curve for LGG (see Appendix D.2, available at [Best vs Worst Prognosis Patients for LGG](#)). For instance, in Fig. 7, the green curve corresponds to a 35-year-old patient with a favorable prognosis, while the red curve represents an 80-year-old patient with a much poorer expected outcome. In addition to age, other relevant features such as prior treatments or molecular characteristics can also be examined, enabling a more comprehensive and transparent clinical interpretation of the model's predictions.

For the patient with the worst prognosis, we observed that in the CNV omics modality, genes such as TP53 appeared amplified. This gene is a well-known prognostic factor in lower-grade glioma⁶⁶. This illustrates how our model enables clinically meaningful interpretation: the predicted survival curve becomes more than just a number when a clinician can relate it to specific omic patterns. Such patient-level predictions can be obtained for any test case, enabling clinicians to provide personalized prognosis.

To further demonstrate the personalized predictions, Fig. 8 shows the survival curves for Patient 0 predicted by SAMVAE using different modality combinations, alongside the KM estimator for the entire cohort. The KM curve represents the average survival trajectory across all patients, while the patient-specific curves illustrate how SAMVAE's predictions for this individual differ from the overall cohort prognosis. The code for generating these results can also be found in⁵⁶.




Fig. 7. Model-predicted survival for patients with best and worst prognoses, alongside the cohort-level KM curve. The mean predicted survival curve is shown with a 90% confidence band (5th–95th percentiles), while the KM curve includes a 95% confidence interval. For the patient with the worst prognosis, clinical and selected omic data were reviewed, revealing an 80-year-old individual with apparent TP53 activation. Such patient-level predictions can be obtained for any test case, enabling clinicians to provide personalized prognosis. An interactive notebook that allows users to explore the clinical features associated with each patient’s survival curve is available at⁵⁶.




Fig. 8. Comparison of survival curves for an individual patient (Patient 0) in LGG-SA predicted by SAMVAE using different combinations of modalities. The predicted curves are shown with 90% confidence intervals (5th–95th percentiles). The KM 95% confidence interval curve represents the average survival trajectory across all patients.

Conclusion

In this work, we introduced SAMVAE, a novel deep learning architecture designed to address the challenges of survival prediction in oncology through effective multimodal data integration while modeling continuous time. By combining clinical information, molecular profiles, and histopathological features into a shared latent space, SAMVAE achieves robust and accurate survival prediction while preserving modality-specific representations. In addition, this unified probabilistic framework enables the generation of personalized and clinically interpretable survival curves.

SAMVAE achieves competitive performance in the multimodal survival analysis setting compared to using clinical data alone, yielding comparable C-index results and consistent improvements in IBS. The superior calibration reflected in the IBS highlights the model’s strength in estimating reliable survival probabilities. This is especially relevant in clinical settings where accurate probabilistic predictions over time are more

informative than mere ranking of patients. Importantly, these gains were achieved using simple and reproducible preprocessing pipelines, suggesting that the benefits of multimodal learning may be even greater with more advanced feature extraction strategies.

In the case of competing risks settings, our results indicate a successful multimodal integration, as SAMVAE achieves competitive performance compared to models trained solely on clinical data. This demonstrates that incorporating omics and histological information contributes meaningful, complementary insights beyond standard clinical variables. In this context, this work presents a multimodal deep learning framework that explicitly addresses competing risks by modeling time to a specific event in continuous domains. Consequently, this framework provides a systematic approach for personalized survival prediction across heterogeneous data sources.

In conclusion, SAMVAE provides a successful framework for multimodal integration, effectively combining omics, imaging, and clinical data into a shared latent space. In addition, our SAMVAE model achieves competitive performance with the current state-of-the-art while relying on a significantly simpler and more unified architecture.

However, it is important to emphasize that SAMVAE is better viewed as a robust fusion model rather than a performance-boosting one. By leveraging multimodal integration to improve calibration (IBS), the model ensures stable risk stratification. This robustness in combining clinical, omics, and imaging data allows for more reliable survival predictions, positioning SAMVAE as a methodological tool with the potential to support risk stratification and multimodal integration in clinical research settings, rather than as a model that directly improves clinical decision-making.

Thanks to its end-to-end design, SAMVAE eliminates the need for complex training pipelines or multiple specialized and pretrained modules. In contrast to black-box models, SAMVAE's probabilistic outputs and interpretable structure facilitate clinical insight, allowing practitioners to assess how different modalities influence risk over time. This interpretability, combined with its ability to handle competing risks, a setting not addressed by existing multimodal survival models, positions SAMVAE as a powerful and practical tool for multimodal survival analysis in biomedical research.

Merits and limitations of the proposed approach

Merits

Personalized and probabilistic survival prediction SAMVAE enables individualized survival prediction by leveraging multimodal data to generate patient-specific survival curves with uncertainty estimates. Unlike traditional survival models that rely on fixed covariate effects, our approach captures complex, nonlinear interactions across clinical and molecular features. This allows us to stratify patients beyond standard clinical risk groups and to identify those who may benefit from intensified follow-up or targeted interventions. As shown in Fig. 7, this capability provides a methodology for future studies in precision medicine, where prognosis could be tailored to the unique characteristics of each patient.

Integration of complementary multimodal information As expected, clinical variables, particularly age, dominate survival prediction, which aligns with existing clinical knowledge and confirms the relevance of classical prognostic factors⁶² (see clinical explainability in Appendices 10 and 11 for further details). However, despite the strong predictive power of clinical variables alone, our study demonstrates that integrating additional omic and imaging modalities leads to well-calibrated models. The consistent reduction in the IBS when combining modalities suggests that non-clinical data contributes meaningful complementary information.

Continuous-time survival modeling The proposed framework models time-to-event outcomes in continuous domains using a Weibull survival distribution, which is widely adopted in survival analysis⁵⁰ and does not rely on proportional hazards assumptions. The approach can also handle discrete-time distributions provided their log-likelihood is differentiable, offering flexibility to adapt the model to different survival modeling settings.

Capability to handle competing risks SAMVAE naturally extends to competing risks settings, enabling the modeling of time to specific events within a unified multimodal framework, a scenario that remains underexplored in multimodal deep survival analysis.

Limitations

Simplified preprocessing and feature engineering A key limitation of our approach is the intentional simplicity in the preprocessing and feature engineering steps. By avoiding extensive preprocessing and complex feature extraction, we highlight the considerable untapped potential of omics and imaging data. While this simplicity ensures that the model remains accessible and reproducible, it may also limit the model's ability to exploit the potential of these datasets fully. However, this deliberate choice guarantees that the entire pipeline is documented and reproducible, ensuring that the methodology can be easily replicated and extended by other researchers. As a result, the observed improvements from modality integration likely represent conservative estimates of the true potential of multimodal fusion.

Limited availability of multimodal benchmarks Given the lack of available multimodal benchmarks in this area that include clinical, omics, and imaging data, direct comparisons with state-of-the-art models are not currently feasible. This absence of benchmarks limits our ability to evaluate SAMVAE's competitive performance fully. Nonetheless, by releasing our code and models publicly, we aim to contribute to the development of future multimodal benchmarks, providing a valuable resource for further comparison and model evaluation in this field.

Cohort size and event imbalance Although the cohort size is relatively modest compared to the complexity of the model, this is a well-recognized challenge in the field of survival analysis, where the availability of large, high-dimensional, and multimodal datasets is often limited (as noted in¹⁰). Additionally, in the context of competing risks, the relatively small number of patients per event type may lead to class imbalance, which can affect the stability of cause-specific survival estimates. Both of these limitations increase the potential for overfitting to dataset-specific patterns. To mitigate these challenges, we employed K-fold cross-validation and reported both the mean and variance of our results, ensuring that SAMVAE's performance remained consistent and was not unduly influenced by any particular subset of data. Nevertheless, we acknowledge that using the same cross-validation framework for both hyperparameter selection and performance evaluation may introduce some degree of performance overestimation in the reported results. While nested cross-validation or an external test set would provide a more stringent assessment of generalization, these were not feasible given the limited size of the available multimodal cohorts.

Limited feature-level interpretability Although the probabilistic formulation provides clinically interpretable survival curves, identifying the specific contribution of individual omics or imaging features to model predictions remains challenging. Further development of explainability techniques will be necessary to link model outputs with clinically actionable knowledge better. In this regard, we have provided additional interpretability analyses in Appendices D and E to further explore the impact of individual features on the survival predictions. These include interactive visualizations of survival curves and partial dependence plots to assess the influence of clinical variables on survival outcomes.

Future directions

Several directions remain open as future work for improving the proposed SAMVAE framework:

- **Modality-aware weighting loss:** In its current form, SAMVAE treats all input modalities uniformly. Introducing modality-specific weighting mechanisms, as suggested in⁵¹, such as those inspired by the weighted Cox hazard approach in³⁷. It could allow the model to dynamically prioritize more informative data sources during training, potentially improving predictive performance and interpretability.
- Although omics and imaging data were preprocessed to ensure basic compatibility, we did not conduct a systematic evaluation of how different preprocessing pipelines affect downstream performance. Future work will investigate the impact of alternative normalization, feature selection, and dimensionality reduction techniques on further enhancing the robustness and generalization of multimodal integration across diverse biomedical datasets.
- Another direction is the development of methods for generating synthetic multimodal patient profiles, which could revolutionize how we address data limitations and resource constraints in clinical practice. A key question for future exploration is whether missing modalities, such as imaging or omics data, can be accurately generated from available modalities (e.g., clinical and genomic data) while maintaining reliable survival predictions. This capability could enable more equitable healthcare by providing robust prognoses for patients with fewer available tests, addressing disparities in resource-limited settings.

Data availability

The data used in this study were obtained from The Cancer Genome Atlas (TCGA), a publicly accessible repository available through the GDC Data Portal: <https://portal.gdc.cancer.gov>. In particular, this work focuses on the BRCA and LGG cohorts available in TCGA. A step-by-step guide detailing the data download process, the applied filters, and the preprocessing pipeline used to construct the BRCA-SA, LGG-SA, BRCA-CR, and LGG-CR datasets is provided in the following public repository: https://github.com/AlbaGarridoLopez/SAMVAE/tree/main/data_download (accessed on 7 July 2025).

Code availability

The code used in this research is publicly available and can be found in the following repository: <https://github.com/AlbaGarridoLopez/SAMVAE> (accessed on 30 January 2026).

Materials availability

Not applicable.

Appendix A Notation

We summarize below the notation used throughout the paper:

General Variables

- N : Total number of samples (patients) in the dataset.
- K : Number of possible event types (e.g., different causes of failure in competing risks).
- L : Number of covariates (features) per sample.
- M : Number of input modalities (e.g., clinical, omics, images).

Data Structure

Each sample is represented by a triplet:

$$(x_i, t_i, y_i) \quad \text{for } i = 1, \dots, N$$

where:

- $x_i^{(m)} \in \mathbb{R}^{L_m}$: Covariate vector of modality m for the i -th individual, where L_m denotes the number of features associated with modality m . Each modality may have a different dimensionality.
- $t_i \in \mathbb{R}_{\geq 0}$: Observed time (event or censoring).
- $y_i \in \{0, 1\}$ for standard survival analysis, or $y_i \in \{0, 1, \dots, K\}$ for competing risks.
- $y_i = 0$: Right-censored observation (no event observed during follow-up).
- $y_i = k \in \{1, \dots, K\}$: Event of type k occurred.

Model Components

- ϕ : Encoder functions for each modality.
- θ : Decoder for covariates.
- ψ : Decoder for survival time(s).
- ω : Classifier for event type (if competing risks).
- $z^{(m)}$: Latent representation of modality m

Notes

- In standard survival analysis, only a single event type is considered ($K = 1$), and $y_i \in \{0, 1\}$.
- In competing risks settings, multiple types of events are modeled, and the event indicator y_i explicitly encodes the type of event or censoring.

Appendix B Parameters tested and optimal

This section summarizes the hyperparameter values explored and the optimal configurations selected for each modality and task (survival analysis and competing risks). Table 4 presents the range of latent dimensions and hidden layer sizes tested during model development. For the omic and imaging modalities, hyperparameter tuning was performed in the presence of clinical data, ensuring that the selected configurations reflect the most effective multimodal integration scenarios. Table 5 reports the optimal parameter combinations for each modality, dataset (LGG and BRCA), and task (standard survival analysis and competing risks modeling).

Mode	Hyperparameter	Values tested
Clinical	Latent dimension	5
	Hidden size	5, 10, 25, 50, 75, 100, 500
Omics	Latent dimension	5, 50, 500
	Hidden size	5, 50, 500
Images	Latent dimension	5, 50, 500
	Hidden size	{8, 16, 32}, {16, 32, 64}, {32, 64, 128}

Table 4. Hyperparameters tested for each modality.

Dataset	Modality	Survival analysis		Competing risks	
		Latent dimension	Hidden size	Latent dimension	Hidden size
LGG	Clinical	10	10	5	75
	DNAm	5	50	5	50
	CNV	5	500	5	500
	miRNA	5	5	5	50
	RNAseq	5	500	5	50
	Patches	5	[8, 16, 32]	5	[16, 32, 64]
BRCA	Clinical	10	100	10	500
	DNAm	5	5	5	5
	CNV	5	5	5	50
	miRNA	5	5	5	5
	RNAseq	5	500	50	50
	Patches	5	[32, 64, 128]	50	[16, 32, 64]

Table 5. Optimal parameters (Latent Dimension and Hidden Size) for each modality, dataset, and scenario.

Appendix C Intermediate results

C.1 Hyperparameter tuning

As an illustrative example, we include in this appendix a subset of the intermediate tables generated during the hyperparameter optimization process for survival analysis and competing risks scenarios. All omics and imaging modalities were optimized in conjunction with clinical data, which served as the baseline for multimodal integration. Additional hyperparameter tuning results and complete tables for other modality combinations and scenarios (BRCA-SA, BRCA-CR, LGG-SA and LGG-CR) are available in the public repository⁵⁶, which also includes the 95% confidence intervals for all reported performance metrics in the supplementary material. Additionally, this repository provides results for all evaluated random seeds to ensure full transparency and reproducibility.

To assess the impact of latent space dimensionality on model performance, we conducted a sensitivity analysis across different latent dimensions. Figure 9 shows the CI-IBS for the top 3 performing configurations in LGG-SA. The results indicate that lower latent dimensions generally achieve better performance compared to higher dimensions, with the exception of the clinical data.

Specifically, we present one of the tables used to select the optimal latent dimension for the LGG-SA prediction task, based on omics data (DNAm) combined with clinical variables (see Table 6). In this table, the first value of the latent dimension and the hidden size correspond to the optimal value previously determined for the clinical data alone and are therefore fixed across all combinations. The second value represents the latent dimension tested for the DNAm modality. Discarding all combinations with any corrected *p*-value below 0.01, the configuration with the best C-index–IBS trade-off among the remaining options is highlighted in bolditalics. In this case, the optimal combination for DNAm corresponds to a latent space dimension of 5 and a hidden size of 50.




Fig. 9. Sensitivity analysis of latent space dimensionality on LGG-SA performance. CI-IBS for the top 3 latent dimension configurations in LGG-SA, showing superior performance with smaller latent spaces.

Latent	Hidden	C-index	IBS	CI - IBS	p-CI HB	p-IBS HB
[10, 5]	[10, 5]	0.685 ± 0.078	0.204 ± 0.027	0.481 ± 0.093	0.318	0.096
[10, 5]	[10, 50]	0.717 ± 0.066	0.183 ± 0.029	0.533 ± 0.086	0.996	0.500
[10, 5]	[10, 500]	0.717 ± 0.051	0.192 ± 0.025	0.525 ± 0.057	0.996	0.404
[10, 50]	[10, 5]	0.610 ± 0.053	0.217 ± 0.028	0.393 ± 0.067	0.000	0.014
[10, 50]	[10, 50]	0.603 ± 0.069	0.223 ± 0.019	0.380 ± 0.067	0.000	0.000
[10, 50]	[10, 500]	0.664 ± 0.061	0.205 ± 0.025	0.459 ± 0.062	0.036	0.096
[10, 500]	[10, 5]	0.556 ± 0.038	0.207 ± 0.019	0.348 ± 0.047	0.000	0.042
[10, 500]	[10, 50]	0.601 ± 0.065	0.217 ± 0.038	0.384 ± 0.082	0.000	0.042
[10, 500]	[10, 500]	0.589 ± 0.073	0.253 ± 0.054	0.337 ± 0.071	0.000	0.000

Table 6. Parameter combination for DNAm with clinical data for LGG-SA. The *p*-values corrected with HB that exceed 0.01 are in bold. Discarding combinations with any *p*-value below 0.05, the one with the best CI-IBS is highlighted in bolditalics. p-CI HB and p-IBS HB represent the *p*-values for the C-Index and IBS corrected with HB, respectively. Detailed results for 95% confidence intervals are available in the supplementary material⁵⁶.

C.2 Intermediate configurations

Additionally, once the optimal hyperparameters were selected for each individual modality (including images), we systematically evaluated all possible combinations of omic modalities and experimented with different numbers of image patches per patient to determine the best configuration for image-based models. As examples, we include the table summarizing the performance of all omic modality combinations for the BRCA-SA task (Table 7), and another showing the image patch selection process for LGG-CR (Table 8). For instance, in the BRCA-SA case, all combinations except one are highlighted in bold, as they surpass the statistical threshold of p -value > 0.01 . Notably, the combination of clinical and CNV modalities achieves the best CI-IBS score among all, and is therefore highlighted in bolditalics to emphasize its superior performance.

In Fig. 8, we show the evaluation of different numbers of patches for BRCA-CR using clinical data. The p -values greater than 0.01 are shown in bold. The configuration with the best CI-IBS (defined as the average of the C-Indices across all risks minus the average IBS) is highlighted in bolditalics. CI-p HB and IBS-p HB represent the p -values for the C-Index and IBS, respectively. In this case, the selected configuration is the one using 10 patches.

After completing the hyperparameter optimization for all modalities and selecting the best combination of omics modalities and the optimal number of image patches, we obtained the final results. These are the outcomes discussed in the Results section (see Sections 5.1 and 5.2). In both intermediate and final experiments, 95% confidence intervals were computed for all performance metrics and are available in the supplementary material provided in the public repository⁵⁶. Additionally, this repository provides results for all evaluated random seeds to ensure full transparency and reproducibility.

Clinical	DNAm	CNV	miRNA	RNAseq	C-index	IBS	CI - IBS	p-CI HB	p-IBS HB
•	•				0.661 ± 0.071	0.200 ± 0.042	0.461 ± 0.086	1.000	0.768
•		•			0.684 ± 0.053	0.175 ± 0.034	0.509 ± 0.069	1.000	1.000
•			•		0.687 ± 0.072	0.207 ± 0.039	0.480 ± 0.084	1.000	0.219
•				•	0.654 ± 0.051	0.186 ± 0.020	0.468 ± 0.043	1.000	1.000
•	•	•			0.641 ± 0.041	0.208 ± 0.039	0.432 ± 0.047	0.350	0.191
•	•		•		0.604 ± 0.050	0.197 ± 0.030	0.407 ± 0.070	0.012	0.628
•	•			•	0.628 ± 0.039	0.186 ± 0.019	0.442 ± 0.046	0.098	1.000
•		•	•		0.599 ± 0.075	0.211 ± 0.046	0.387 ± 0.086	0.027	0.193
•		•		•	0.617 ± 0.061	0.195 ± 0.025	0.423 ± 0.064	0.076	0.760
•			•	•	0.629 ± 0.073	0.193 ± 0.028	0.436 ± 0.086	0.321	1.000
•	•	•	•		0.624 ± 0.060	0.196 ± 0.029	0.428 ± 0.069	0.135	0.665
•	•	•		•	0.587 ± 0.042	0.185 ± 0.021	0.403 ± 0.048	0.001	1.000
•	•		•	•	0.608 ± 0.053	0.192 ± 0.027	0.417 ± 0.071	0.021	1.000
•		•	•	•	0.609 ± 0.046	0.194 ± 0.023	0.415 ± 0.043	0.017	0.770
•	•	•	•	•	0.617 ± 0.058	0.183 ± 0.025	0.434 ± 0.056	0.065	1.000

Table 7. Combination of clinical data with pmic Scenarios for BRCA-SA with a single risk. The p -values that exceed 0.01 are in bold. The one with the best CI-IBS is highlighted in bolditalics. CI-p HB and IBS-p HB represent the p -values corrected with HB for the C-Index and IBS, respectively. Detailed results for 95% confidence intervals are available in the supplementary material⁵⁶.

N Patches	Risk	C-index	IBS	AVG CI-IBS	p-CI HB	p-IBS HB
1	1	0.583 ± 0.034	0.231 ± 0.051	0.369	0.011	1.000
1	2	0.622 ± 0.042	0.236 ± 0.057	0.369	1.000	0.915
5	1	0.607 ± 0.049	0.243 ± 0.049	0.387	0.891	0.616
5	2	0.625 ± 0.074	0.215 ± 0.031	0.387	1.000	1.000
10	1	0.627 ± 0.036	0.217 ± 0.027	0.400	1.000	1.000
10	2	0.614 ± 0.054	0.224 ± 0.043	0.400	1.000	1.000
15	1	0.601 ± 0.057	0.220 ± 0.030	0.380	0.675	1.000
15	2	0.612 ± 0.052	0.233 ± 0.060	0.380	1.000	1.000

Table 8. Evaluation of different numbers of patches for BRCA-CR. The p -values greater than 0.01 are in bold. The configuration with the best CI-IBS (defined as the average of the C-Indices across all risks minus the average IBS) is highlighted in bolditalics. CI-p HB and IBS-p HB represent the p -values corrected with HB for the C-Index and IBS, respectively. Detailed results for 95% confidence intervals are available in the supplementary material⁵⁶.

Appendix D Multimedia materials

This section contains interactive visualizations and exploratory tools that complement the core results of our study. These materials allow dynamic inspection of survival predictions across models and individual patients, as well as gene-specific queries for personalized genomic insights.

D.1 Interactive survival curves by model

An interactive visualization comparing the survival curves of different trained models against the KM estimate. Users can toggle models on/off via the legend and visually compare model behaviors across different data modalities (clinical, omic, histological). The plot also enables inspection of the estimated survival probability at each time point for every model.

- [Interactive LGG Survival Plot](#)
- [Interactive BRCA Survival Plot](#)




Fig. 10. Beeswarm plot of SHAP values for LGG-SA.




Fig. 11. Mean Absolute SHAP Values for LGG-SA Clinical Features.

D.2 Best versus worst prognosis patients

A personalized medicine tool that visualizes survival predictions for the best and worst prognosis patients according to the model. For each patient, the plot shows the mean predicted survival curve with 90% confidence interval and the clinical variables as hover-over tooltips. This tool helps interpret the survival curves derived from clinical data and demonstrates that it is possible to generate personalized survival curves for each patient, thereby promoting personalized medicine and supporting model explainability:

- [Best vs Worst Prognosis Patients for LGG](#)

Additionally, a Gene CNV Query Notebook is provided. This notebook enables users to check the CNV status of frequently mutated genes for each patient. Beyond visualizing survival curves, users can verify the presence and alteration of specific genes known to be cancer risk indicators on a per-patient basis.

Appendix E Clinical explainability

To improve interpretability, SHAP (SHapley Additive exPlanations)⁶⁷ values were computed for the clinical modality. These values quantify the contribution of each clinical feature to individual survival predictions. Beeswarm and bar plots (Figs. 11 and 10) were generated to visualize both the global importance and local effects of features.

In Figs. 11 and 10, we can see that age-related features play a dominant role in survival prediction. The model consistently assigns higher importance to “age at diagnosis” and “age at index”, reaffirming that patient age is a key factor in determining prognosis, as previously discussed in⁶².

To further assess the relationship between clinical variables and survival outcomes, we calculated partial dependence of survival functions on selected clinical variables that were identified as important in the SHAP analysis, including age and gender, and the CNV status of the TP53 gene, because it is a well-known prognostic factor in lower-grade glioma⁶⁶. These stratified survival curves demonstrate the individual prognostic impact of each variable. The log-rank test⁶⁸ was applied to compute p-values, identifying significant prognostic factors. To avoid redundancy, we only present Fig. 12. The corresponding code is available in⁵⁶ along with the remaining figures.

In LGG-SA (Fig. 12), age was found to be highly significant (p -value < 0.05), highlighting its strong influence on survival outcomes. In contrast, gender did not show a significant impact on prognosis (p -value > 0.05). Notably, when TP53 is amplified, it is associated with a worse prognosis, which aligns with existing literature on the role of TP53 mutations in glioma survival⁶⁶.




Fig. 12. Partial dependence plots showing the relationship between survival outcomes and selected clinical variables (age, gender, and the CNV status of the TP53 gene) in the LGG-SA dataset, with shaded regions indicating 95% confidence intervals. p-values were computed using the log-rank test to assess statistical significance.

Received: 10 December 2025; Accepted: 9 June 2026

Published online: 18 June 2026

References

- Salvi, M. et al. Multi-modality approaches for medical support systems: A systematic review of the last decade. *Inf. Fusion* **103**, 102134. <https://doi.org/10.1016/j.inffus.2023.102134> (2024).
- Waqas, A., Tripathi, A., Ramachandran, R. P., Stewart, P. A. & Rasool, G. Multimodal data integration for oncology in the era of deep neural networks: a review. *Front. Artif. Intell.* **7**, 1408843 (2024).
- Boehm, K. M., Khosravi, P., Vanguri, R., Gao, J. & Shah, S. P. Harnessing multimodal data integration to advance precision oncology. *Nat. Rev. Cancer* **22**(2), 114–126 (2022).
- Gutman, D. A. et al. Cancer digital slide archive: an informatics resource to support integrated in silico analysis of TCGA pathology data. *J. Am. Med. Inf. Associat.* **20**, 1091–1098 (2013) (Epub 2013 Jul 25).
- Meng, C. et al. Dimension reduction techniques for the integrative analysis of multi-omics data. *Brief. Bioinf.* **17**(4), 628–641. <https://doi.org/10.1093/bib/bbv108> (2016).
- Wang, F., Oh, T. W., Vergara-Niedermayr, C., Kurc, T. & Saltz, J. Managing and querying whole slide images. *Procd. SPIE Int. Soc. Opt. Eng.* **8319**, 83190. <https://doi.org/10.1117/12.912388> (2012).
- Olivier, M., Asmis, R., Hawkins, G. A., Howard, T. D. & Cox, L. A. The need for multiomics biomarker signatures in precision medicine. *Int. J. Mole. Sci.* **20**, 4781 (2019).
- Altman, D., De Stavola, B., Love, S. & Stepniowska, K. Review of survival analyses published in cancer journals. *British J. Cancer* **72**(2), 511–518 (1995).
- Prentice, R. L. et al. The analysis of failure times in the presence of competing risks. *Biometrics* **32**, 541–554 (1978).
- Wiegrebe, S., Kopper, P., Sonabend, R., Bischl, B. & Bender, A. Deep learning for survival analysis: a review. *Artif. Intell. Rev.* **57**(3), 65 (2024).
- Kingma, D. P. et al. *Auto-encoding variational bayes* (Banff, Canada, 2013).
- Apellániz, P. A., Parras, J. & Zazo, S. Leveraging the variational Bayes autoencoder for survival analysis. *Sci. Rep.* **14**(1), 24567 (2024).
- Cox, D. R. Regression models and life-tables. *J. Royal Statist. Soc. Ser. B (Methodological)* **34**(2), 187–202 (1972).
- Ravi, D. et al. Deep learning for health informatics. *IEEE J. Biomed. Health Inf.* **21**(1), 4–21 (2016).
- Denfeld, Q. E., Burger, D. & Lee, C. S. Survival analysis 101: An easy start guide to analysing time-to-event data. *Eur. J. Cardiovascular Nurs.* **22**, 332–337 (2023).
- Apellániz, P. A., Parras, J. & Zazo, S. CR-SAVAE: A Parametric Method for Survival Analysis with Competing Risks 1526–1530 (2024). <https://doi.org/10.23919/EUSIPCO63174.2024.10715431>
- Katzman, J. L. et al. DeepSurv: Personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Med. Res. Methodol.* **18**, 1–12. <https://doi.org/10.1186/s12874-018-0482-1> (2018).
- Kvamme, H., Borgan, Ø. & Scheel, I. Time-to-event prediction with neural networks and Cox regression. *J. Mach. Learn. Res.* **20**(129), 1–30 (2019).
- Hao, J., Kim, Y., Mallavarapu, T., Oh, J. H. & Kang, M. K. Cox-PASNet: Pathway-based sparse deep neural network for survival analysis 381–386 (2018). <https://doi.org/10.1109/BIBM.2018.8621262>.
- Lee, C., Zame, W., Yoon, J., & Schaar, M.V.D. DeepHit: a deep learning approach to survival analysis with competing risks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32 (2018). <https://ojs.aaai.org/index.php/AAAI/article/view/11858>
- Lee, C., Yoon, J. & Schaar, M. V. D. Dynamic-DeepHit: A deep learning approach for dynamic survival analysis with competing risks based on longitudinal data. *IEEE Trans. Biomed. Eng.* **67**(1), 122–133. <https://doi.org/10.1109/TBME.2019.2912173> (2019).
- Bennis, A., Mouysset, S. & Serrurier, M. Estimation of Conditional Mixture Weibull Distribution with Right Censored Data Using Neural Network for Time-to-Event Analysis. In *Advances in Knowledge Discovery and Data Mining: 24th Pacific-Asia Conference* 687–698 (Springer, Singapore, 2020).
- Nagpal, C., Li, X. & Dubrawski, A. Deep survival machines: Fully parametric survival regression and representation learning for censored data with competing risks. *IEEE J. Biomed. Health Inf.* **25**(8), 3163–3175. <https://doi.org/10.1109/JBHI.2021.3073206> (2021).
- Fine, J. P. & Gray, R. J. A proportional hazards model for the subdistribution of a competing risk. *J. Am. Statist. Associat.* **94**(446), 496–509 (1999).
- Li, Y., Jia, W., Kang, Y., Chen, T., Li, X., Du, X., Dong, J., Ma, C., Wang, F., & Xie, G. Deepcomp: Which competing event will hit the patient first? In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* 629–636 (IEEE, 2020).
- Gupta, G., Sunder, V., Prasad, R., & Shroff, G. CRESA: a deep learning approach to competing risks, recurrent event survival analysis. In *Advances in Knowledge Discovery and Data Mining: 23rd Pacific-Asia Conference, PAKDD 2019, Macau, China, April 14–17, 2019, Proceedings, Part II* 23 108–122 (Springer, 2019).
- Chi, S. et al. Deep semisupervised multitask learning model and its interpretability for survival analysis. *IEEE J. Biomed. Health Inform.* **25**(8), 3185–3196 (2021).
- Huang, P., Liu, Y., et al. DeepCompete: A deep learning approach to competing risks in continuous time domain. In *AMIA Annual Symposium Proceedings*, vol. 2020, 177 (2021)
- Thorsen-Meyer, H.-C. et al. Discrete-time survival analysis in the critically ill: a deep learning approach using heterogeneous data. *NPJ Digit. Med.* **5**(1), 142. <https://doi.org/10.1038/s41746-022-00679-6> (2022).
- Tang, B., Li, A., Li, B. & Wang, M. CapSurv: capsule network for survival analysis with whole slide pathological images. *IEEE Access* **7**, 26022–26030. <https://doi.org/10.1109/ACCESS.2019.2901049> (2019).
- Giordano, T. J. The cancer genome atlas research network: a sight to behold. *Endocr. Pathol.* **25**, 362–365 (2014).
- Liu, P., Ji, L., Ye, F. & Fu, B. AdvMIL: Adversarial multiple instance learning for the survival analysis on whole-slide images. *Med. Image Anal.* **91**, 103020. <https://doi.org/10.1016/j.media.2023.103020> (2024).
- Hao, J., Kosaraju, S.C., Tsaku, N.Z., Song, D.H., & Kang, M. PAGE-Net: Interpretable and Integrative Deep Learning for Survival Analysis Using Histopathological Images and Genomic Data 355–366. https://doi.org/10.1142/9789811215636_0032 . https://www.worldscientific.com/doi/abs/10.1142/9789811215636_0032
- Tong, L., Mitchel, J., Chatlin, K. & Wang, M. D. Deep learning based feature-level integration of multi-omics data for breast cancer patients survival analysis. *BMC Med. Inform. Decis. Mak.* **20**(1), 225. <https://doi.org/10.1186/s12911-020-01225-8> (2020).
- Lin, J. et al. GD-Net: an integrated multimodal information model based on deep learning for cancer outcome prediction and informative feature selection. *J. Cell Mol. Med.* **28**(23), 70221 (2024).
- Vale-Silva, L. A. & Rohr, K. Long-term cancer survival prediction using multimodal deep learning. *Sci. Rep.* **11**(1), 13505. <https://doi.org/10.1038/s41598-021-92799-4> (2021).
- Mondol, R. K., Millar, E. K. A., Sowmya, A. & Meijering, E. BioFusionNet: deep learning-based survival risk stratification in ER+ breast cancer through multifeature and multimodal data fusion. *IEEE J. Biomed. Health Inform.* **28**(9), 5290–5302. <https://doi.org/10.1109/jbhi.2024.3418341> (2024).
- Wei, T. et al. Survival prediction of stomach cancer using expression data and deep learning models with histopathological images. *Cancer Sci.* **114**(2), 690–701 (2023).

39. Chen, R.J., Lu, M.Y., Weng, W.-H., Chen, T.Y., Williamson, D.F., Manz, T., Shady, M., & Mahmood, F. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* 4015–4025 (2021).
40. Li, R., Wu, X., Li, A. & Wang, M. HFBSurv: hierarchical multimodal fusion with factorized bilinear models for cancer survival prediction. *Bioinformatics* **38**(9), 2587–2594 (2022).
41. Chen, R. J. et al. Pathomic fusion: an integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. *IEEE Trans. Med. Imaging* **41**(4), 757–770 (2020).
42. Lv, Z., Lin, Y., Yan, R., Wang, Y. & Zhang, F. TransSurv: transformer-based survival analysis model integrating histopathological images and genomic data for colorectal cancer. *IEEE/ACM Trans. Comput. Biol. Bioinf.* **20**(6), 3411–3420 (2022).
43. Kopper, P., Wiegerebe, S., Bischl, B., Bender, A. & Rügamer, D. DeepPamm: Deep piecewise exponential additive mixed models for complex hazard structures in survival analysis. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining* 249–261 (Springer, 2022).
44. Jeanselme, V., Yoon, C.H., Tom, B., & Barrett, J. Neural Fine-Gray: Monotonic neural networks for competing risks. In: Mortazavi, B.J., Sarker, T., Beam, A., Ho, J.C. (eds.) *Proceedings of the Conference on Health, Inference, and Learning. Proceedings of Machine Learning Research* vol. 209, 379–392. (PMLR, 2023). <https://proceedings.mlr.press/v209/jeanselme23a.html>
45. Groha, S., Schmon, S. M. & Gusev, A. A general framework for survival analysis and multi-state modelling (2020) [arXiv:2006.04893](https://arxiv.org/abs/2006.04893) arXiv preprint.
46. Kim, S., Kazmierski, M. & Haibe-Kains, B. Deep-CR MTLR: A multi-modal approach for cancer survival prediction with competing risks. In *Survival Prediction- Algorithms, Challenges and Applications* 223–231 (PMLR, 2021).
47. Kaplan, E. L. & Meier, P. Nonparametric estimation from incomplete observations. *J. Am. Stat. Assoc.* **53**(282), 457–481 (1958).
48. Kalbfleisch, J. D. & Prentice, R. L. Marginal likelihoods based on Cox's regression and life model. *Biometrika* **60**(2), 267–278 (1973).
49. Goldblum, M., Finzi, M., Rowan, K. & Wilson, A. G. *The No Free Lunch Theorem, Kolmogorov Complexity, and the Role of Inductive Biases in Machine Learning* (2024) [arXiv:org/abs/2304.05366](https://arxiv.org/abs/2304.05366).
50. Ranganath, R., Perotte, A., Elhadad, N. & Blei, D. Deep survival analysis. In *Machine Learning for Healthcare Conference* 101–114 (PMLR, 2016).
51. Martínez-García, M. & Olmos, P. Handling Ill-conditioned omics data with deep probabilistic models. *IEEE J. Biomed. Health Inf.* <https://doi.org/10.1109/JBHI.2023.3279493> (2023).
52. Liu, J. et al. An integrated TCGA pan-cancer clinical data resource to drive high-quality survival outcome analytics. *Cell* **173**(2), 400–416 (2018).
53. Tong, L., Mitchel, J., Chatlin, K. & Wang, M. D. Deep learning based feature-level integration of multi-omics data for breast cancer patients survival analysis. *BMC Med. Inform. Decis. Mak.* **20**, 1–12 (2020).
54. Vale-Silva, L. A. & Rohr, K. Long-term cancer survival prediction using multimodal deep learning. *Sci. Rep.* **11**(1), 13505 (2021).
55. Lu, M. Y. et al. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nat. Biomed. Eng.* **5**(6), 555–570 (2021).
56. Garrido, A.: SAMVAE: Code for Deep Survival Analysis in Multimodal Medical Data. Accessed: 29 January 2026. <https://github.com/AlbaGarridoLopez/SAMVAE>
57. Antolini, L., Boracchi, P. & Biganzoli, E. A time-dependent discrimination index for survival data. *Stat. Med.* **24**(24), 3927–3944 (2005).
58. Robins, J.M., et al. Information recovery and bias adjustment in proportional hazards regression analysis of randomized trials using surrogate markers. In *Proceedings of the Biopharmaceutical Section*, vol. 24, 3 (American Statistical Association, San Francisco CA, 1993).
59. Henderson, P. et al. Deep reinforcement learning that matters. In *Proceedings of the AAAI Conference on Artificial Intelligence* vol. 32 (2018).
60. Holm, S. A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* **6**, 65–70 (1979).
61. Giacalone, M., Agata, Z., Cozzucoli, P. C. & Alibrandi, A. Bonferroni-holm and permutation tests to compare health data: methodological and applicative issues. *BMC Med. Res. Methodol.* **18**(1), 81 (2018).
62. Schellhout, L. J. et al. Multivariate analysis of survival in differentiated thyroid cancer: the prognostic significance of the age factor. *Eur. J. Cancer Clin. Oncol.* **24**(2), 331–337 (1988).
63. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**(3), 261–272 (2020).
64. Aalen, O. O. & Johansen, S. An empirical transition matrix for non-homogeneous Markov chains based on censored observations. *Scand. J. Stat.* **5**, 141–150 (1978).
65. Therneau, T., Crowson, C., & Atkinson, E. Multi-state models and competing risks. CRAN-R (<https://cran.r-project.org/web/packages/survival/vignettes/compete.pdf>) (2020)
66. Noor, H., Briggs, N. E., McDonald, K. L., Holst, J. & Vittorio, O. TP53 mutation is a prognostic factor in lower grade glioma and may influence chemotherapy efficacy. *Cancers* **13**(21), 5362 (2021).
67. Lundberg, S. M. & Lee, S. -I. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).
68. Bland, J. M. & Altman, D. G. The logrank test. *BMJ* **328**(7447), 1073 (2004).

Author contributions

All authors contributed to the conceptualization of the study. A.G., A.A., and P.A.A. designed the general architecture. A.G. and P.A.A. were responsible for software implementation and validation. Data curation was performed by A.G. Formal analysis was conducted by J.P. and S.Z., who also oversaw project administration and supervision together with A.A. A.G. wrote the original draft of the manuscript. All authors reviewed and approved the final version of the manuscript.

Funding

SYNTHIA (Synthetic Data Generation framework for integrated validation of use cases and AI healthcare applications) is supported by the Innovative Health Initiative Joint Undertaking (IHI JU) under grant agreement No 101172872. The JU receives support from the European Union's Horizon Europe research and innovation programme, COCIR, EFPIA, Europa Bio, MedTech Europe, Vaccines Europe, and DNV. The UK consortium partner, The National Institute for Health and Care Excellence (NICE) is supported by UKRI Grant 10132181. Funded by the European Union, the private members, and those contributing partners of the IHI JU. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the aforementioned parties. Neither of the aforementioned parties can be held responsible for them.

Declarations

Competing interests

The authors declare no competing interests.

Ethics approval and consent to participate

According to the TCGA data usage policies, all data used are publicly available and fully anonymized. Therefore, no ethical approval was required for their use in this research.

Consent for publication

Not applicable.

Additional information

Correspondence and requests for materials should be addressed to A.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026