SSRG International Journal of Electronics and Communication Engineering
ISSN: 2348-8549/ https://doi.org/10.14445/23488549/IJECE-V13I7P103

Volume 13 Issue 7, 35-50, July 2026

Original Article

Multi-Modal Treatment Response Prediction and Recurrence Risk Analysis for Bone Tumors Using A Hybrid Longitudinal Deep Reinforcement Model

Rathla Roopsingh[1], D. Vasumathi[2]

[1,2]*Department of CSE, Jawaharlal Nehru Technological University, Hyderabad, Telangana, India.*

[1]*Corresponding Author : rathlaroopsingh@gmail.com*

Received: 18 April 2026 Revised: 06 June 2026 Accepted: 17 June 2026 Published: 29 July 2026

Abstract *- Predicting treatment response and recurrence risk in bone tumor patients is crucial to the choice of effective therapeutic strategies and the enhancement of clinical outcome. However, currently available prognosis and decision-support tools are largely based on fixed imaging or univariate clinical measures, and are unable to learn over time disease dynamics and therapies' interaction. It is based on progress in tumor segmentation and grading and presents a new Hybrid Longitudinal Deep Reinforcement Model (HL-DRM) that was specifically created to predict multi-modal treatment responses and predict recurrence in bone tumors. The suggested model combines sequential CT/MRI images, radiomic features, and patient clinical history with a Temporal Fusion Transformer to learn patterns of tumor progression over time. A policy agent of reinforcement-learning predicts optimal treatment pathways (surgery, radiotherapy, chemotherapy, or combinations of those) by maximizing the survival reward signals. Moreover, a Recurrence-Aware Survival Module based on Cox neutral modelling is used to predict personalized recurrence probabilities and survival curves. Experimental analysis of a longitudinal bone tumor data set shows that HL-DRM is far more effective at predicting prognosis, response, and estimating risk of recurrence, as compared to standard prognosis models. This direction promotes the AI-based precision oncology and helps clinicians make evidence-based, personalized treatment choices.*

Keywords *- Multi-modal deep learning, Longitudinal medical imaging, Treatment response prediction, Recurrence risk analysis, Reinforcement learning, Temporal Fusion Transformer.*

1. Introduction

Both benign and malignant bone tumors require adequate assessment of treatment outcome and the probability of recurrence over the long term, with the view to inform the clinical rationale and increase the survival rates of the patients [1]. Although the latest imaging systems and Deep Learning approach have enhanced the tumor segmentation and grading process, it is challenging to determine whether a patient will respond to the treatment since tumors are heterogeneous with complex biological behaviour and varying treatment plan per patient [2]. The traditional prognosis methods, such as Cox regression, logistic analysis, or even static radiomics, are not able to consider the time evolutionary patterns and multi-modal patient information, which results in ineffective generalization and clinical accuracy [3].

The latest advances in AI have enabled the application of multi-modal fusion and deep temporal learning, but their application in the modelling of treatment response to bone tumours is not explored sufficiently [4]. Moreover, the existing frameworks lack reinforcement learning to recommend the most optimal course of action in therapy based on the real outcomes of the treatment [5, 6]. To address these limitations, the present work introduces the HL-DRM - a single workflow, comprising of sequential CT/MRI tumor [7] progression analysis [8], radiomic texture-shape-intensity embeddings [9], and structured clinical data [10]. The model represents longitudinal aspects using a Temporal Fusion Transformer and suggests adaptive treatment plans that optimize the predicted survival benefit by a category of reinforcement-learning-based agent. A Recurrence-Aware Survival Module also estimates the recurrence risk curves and patient-specific survival time.

Deep learning has been shown to be very effective at the detection of bone tumors, segmentation, and classification; however, most previous research has been limited to a single task, and few have been developing a full clinical decision support system. There is a paucity of methods that combine longitudinal tumour progression, prediction of responsiveness to treatment, risk of recurrence, and survival analysis. Furthermore, there are a number of existing techniques that are based on static images and do not depict the temporal evolution of tumor features. Recent survival models based on transformers can capture complex relationships between features, but lack Reinforcement Learning mechanisms for adapting treatment optimization. Thus, it is necessary to develop a unified framework to perform accurate analysis of longitudinal multi-modal data,

make predictions of successful treatment, estimate the likelihood of recurrence, and provide customized treatment recommendations.

The Novelty of HL-DRM is a state-of-the-art, all-in-one predictive treatment response and recurrence risk learning architecture in bone tumors that learns longitudinal progression of sequential CT scan and MRI measurements. It uses Temporal Fusion Transformer to fuse dynamic multi-modal features with reinforcement learning to optimize personalized treatment and recurrence-sensitive survival modelling to give a precise individualized prognosis.

The Research contributions of this work are as follows: Development of HL-DRM, a new Hybrid Longitudinal Deep Reinforcement Model, which is specially designed to predict bone tumor response after therapy and assess recurrence risk.

- A multi-modal combination of consecutive CT/MRI images, radiomic profiles, and clinical variables, which allows characterizing tumor dynamics in detail throughout time.
- Progression modelling with Temporal Fusion Transformer to learn the dynamics of tumor behaviour patterns and therapy effects.
- Optimization of treatment strategies through reinforcement-learning, which gives adaptive decision guidance to maximize the reward of survival.
- Recurrence-aware survival estimation with Cox-based neural encoding to predict survival curves individually.
- Wide experimental validation of longitudinal bone tumor data, with better accuracy and robustness compared to traditional prognosis and machine-learning-based survival models.

2. Related Works

This situation with bone tumor diagnosis in children is problematic, as it is not common, and the condition has a general clinical picture. In study [11], the authors studied models of Vision Transformer to classify the healthy and pathological knee X-ray images in children using transfer learning and large-scale augmentation to break the lack of data. It was discovered that the model was very diagnostic and, with explainable AI using Grad-CAM, could help in the early referral decision, which also shows that Deep Learning can be applied in the detection of bone tumors at an early stage. Quantitative evaluation of the bone quality in an oncological setting has also been quantitatively evaluated using Deep Learning. In [12], a 3DResUNet-based architecture was developed to predict the volume bone mineral density of spinal tumor metastatic tumor patients using the QCT scans. The proposed method was closely related to the ground-truth measures and very good diagnostics in the screening of osteoporosis, which indicates that the automated volumetric bone analysis can be done in the metastatic disease setting. The learning approaches based on multitasking have delivered on the assessment of bone tumor in general. Paper [13] developed a single Deep Learning model that can potentially localize bounding boxes, segment, and classify primary bone tumors on radiographs in a single step. The model has been experimented with both internal and external datasets and has been shown to be reliable in a wide range of tasks, which is why it is better to use integrated diagnostic pipelines rather than one-task solutions. The issue of differentiating benign and malignant bone lesions with the aid of MRI has been taken into account in [14], where Efficient Net-based image models have been applied, and the demographic data about the patient were utilized in the form of a voting ensemble. Not only did the system achieve expert-level performance, but it was highly generalized across institutions and enables the application of multimodal Deep Learning to minimize unnecessary biopsies and specialist visits. Radiomics Deep learning has been applied to quantify the extent of the tumor burden in scintigraphy of Bone. In [15], a Bone Scan Index system was automated and tested on different types of cancer organizations by utilizing neural networks. Even though high sensitivity and specificity were seen in prostate, breast, and colorectal cancer, low performance in melanoma and lung cancer showed the variability of bone metastasis in various tumors.

A thorough review of the implementation of Deep Learning in the issue of bone tumors was presented in [16], including the advances in the sphere of detection, segmentation, classification, grading, and prediction of the prognosis basing on radiological and pathological images. The review identified significant barriers, including the limited datasets, ineffectiveness in standardization, and poor external validation, emphasizing the need to have highly effective and generalizable AI models. Large-scale multi-institutional studies have also demonstrated the use of Deep Learning on primary bone tumor classification. In Study [17], the binary and multi-class tumor classification was performed with the use of conventional radiographs of more than 1,300 patients. The model demonstrated the same performance as the subspecialist radiologists, who had an excellent performance in comparison with junior clinicians, and was highly generalizable across institutions. On a region-specific basis, classification of bone tumors has also been explored.

In [18], a Dense Net model was developed to perform proximal tumor classification of plain radiographs of the femur. It has shown good results in both terms of AUROC and better performance in the system compared to clinicians, which means that deep learning is relevant in the anatomically localized bone tumor diagnostics process. Malignant bone lesion segmentation has been widely studied, as automated segmentation of malignant bone lesions has been of interest to clinical users. In [19], the deep learning-based segmentation method in CT, MRI, and PET/CT was systematically reviewed. The findings confirmed the stability of performance increase with the involvement of the variants of U-Net, data augmentation, and pre-processing, and demonstrated the heterogeneity of datasets as one of the serious limitations.

Modern medical foundation models and transformer-based health care analytics have shown robust performance

in capturing long-range dependencies between longitudinal medical records. Medical foundation models and healthcare analytics based on transformer networks have recently emerged as powerful techniques to capture long-range dependencies in longitudinal patient records. Currently, they still have limited applications for bone tumor prognosis, as most studies emphasize the diagnosis and classification of the tumor and not the modelling of the response to treatment. Moreover, other cancer types have already been successfully applied to an adaptive treatment plan for Reinforcement Learning, yet the application of it in the temporal imaging progression analysis in bone tumors is still less explored. This inspires the construction of a single framework that would allow the simultaneous modelling of disease evolve, treatment optimization, and estimation of disease recurrence risk.

Traditional image processing, along with Deep Learning, has been explored in detecting bone cancer. In [20], the CT images were segmented using clustering and edge detection algorithms, and then they were categorized on the CNN architectures. The accuracy of classification was good, suggesting that hybrid techniques can be successfully used in the differentiation of normal and cancerous bone tissue. Deep learning has additionally been applicable in the detection of bone metastasis using nuclear medicine imaging. In one study [21], a CNN-based bone scan classifier of the whole body of breast cancer patients was designed. The suggested model also outperformed the classical CNN models, which once again justifies the helpfulness of Deep Learning in the screening of metastasis.

Besides tumor detection, Deep Learning has been applicable in fracture risk and bone density. In [22], a system was developed that approximated the bone mineral density and fracture risk through plain radiographs by utilizing an automated system that gave high accuracy of calibration and prediction. The device demonstrated that it was possible to screen osteoporosis more effectively through other methods other than special imaging modalities. Osteoporosis screening was also researched in [23], with CT-based screening, where DenseNet models were trained to classify bone density and predict quantitative values in case of automatic segmentation. Independent tests of multiple hospitals checking generalization were crudely checked, proving the power of deep learning in bone health assessment.

Lesion levels of metastatic bone disease have been classified using the state of the art CNN architectures. Four various 3D and 2D deep learning models were experimented in the study [24] using a large and properly labelled CT dataset of prostate cancer patients. The ensemble method could achieve high-quality accuracy, which is an indicator of the importance of the texture, morphology, and volumetric characteristics of the lesion in characterization. Finally, [25] summarized epidemiological knowledge on bone tumours, and they provided a systematic review of the incidence, survival rates, and risk factors of benign and malignant bone tumours. The article emphasized the increasing understanding of bone tumors and its impact on the survival rates in case of their early detection, which is why AI-based diagnostics and prognostics have a clinical rationale.

Table 1. Comparison of prognostic capabilities across survival models

Method	Multi-modal Fusion	Longitudinal Analysis	RL-based Treatment Optimization	Recurrence Prediction	Survival Modelling
DeepSurv	✗	✗	✗	✓	✓
Temporal CNN	✓	✓	✗	✓	✓
Transformer Survival	✓	✓	✗	✓	✓
2D-3D CNN Ensemble	✓	✗	✗	✗	✗
Proposed HL-DRM	✓	✓	✓	✓	✓

Existing solutions are generally based on either modelling survival or characterisation of lesions, as can be seen in the comparison. No single approach combines Longitudinal Progression Learning, Reinforcement-Based Therapeutic Optimization, Recurrence Prediction Based on Survival Analysis, and Multi-modal Fusion together. HL-DRM overcomes these limitations by providing a common framework.

3. Proposed Model

The proposed HL-DRM integrates the multi-modal data of medical information to identify the response and recurrence of bone tumor treatment. Radiomic features and clinical records are then run and integrated with longitudinal CT/MRI images, STSN segmentation outputs, and grading data of RMD-Net outputs. A Temporal Fusion Transformer learns time-varying patterns of tumor progression, producing progression-aware embeddings. They are used to generate a model of the patient condition that can be tracked by a Reinforcement Learning agent to propose the most desirable treatment plans with the help of policy-gradient optimization. At the same time, a Recurrence-Aware Survival Module, which uses Cox-based neural modelling, is used to estimate the recurrence probability and the individual survival curve. The resultant system offers predictions of the treatment response, recurrence risk, and

suggestions of therapy that may be utilized to assist in steering the accuracy of clinical decision-making in bone oncology. Figure 1 illustrates the whole flow of the proposed HL-DRM, starting with multi-modes of input data and tumor-centric features extraction to the temporal progression modelling. Its primary component is the Temporal Fusion Transformer that enables progression-based state building. Reinforcement-based treatment optimization and recurrence-aware survival model co-construct the treatment response forecasting, recurrence risk forecasting, and clinical decision support.

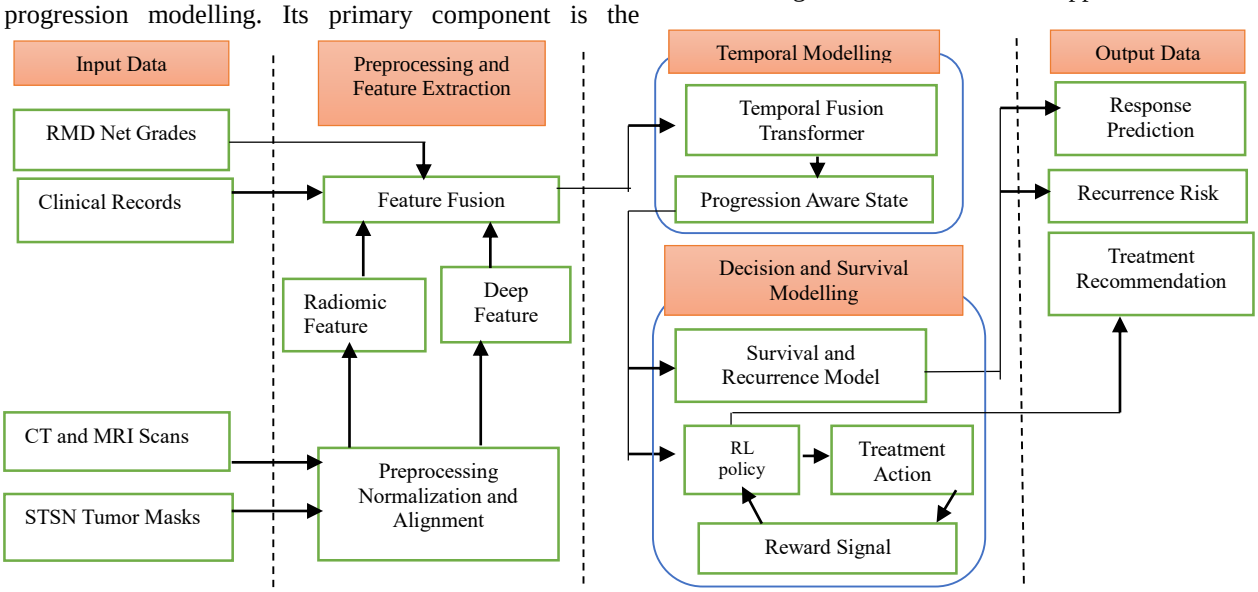

Fig. 1 Overall architecture of the hybrid longitudinal deep reinforcement model

3.1. Multi-Modal Data Acquisition and Preprocessing

The proposed HL-DRM methodology involves a multi-modal data set, which is a longitudinal CT scan and MRI scan sequence in a few clinical follow-ups of bone tumor patients. The Stochastic Tumor Segmentation Network (STSN) and tumor grading outputs generated by RMD-Net are associated with the tumor segmentation mask of every imaging volume. In addition, structured clinical data such as demographic, tumor stage, biochemical measurements, and treatment history (surgery, chemotherapy, and radiotherapy dose cycles) are included to reflect the biological differences as a result of treatment over time. It is possible to describe the time-consistent and enriched progression of the disease of individual patients using such multi-source data. The raw imaging datasets are standardized before additional processing to reduce variations caused by different image acquisition conditions. Normalization of intensities is carried out through min-max scaling, whereby there is a stable distribution of inputs across imaging modalities:

$$I_{norm}(x) = \frac{I(x) - \min(I)}{\max(I) - \min(I)} \qquad (1)$$

Where $I(x)$is the original intensity at voxel x, and $\min(I), \max(I)$denote the global minimum and maximum intensity values. Subsequently, spatial alignment between sequential scans is achieved through affine registration, minimizing positional variations between acquisition sessions. The optimal transformation matrix $T^*$is estimated by maximizing the Normalized Mutual Information (NMI) objective:

$$T^* = \arg\max_T \, NMI(I_{ref}, T(I_t)) \qquad (2)$$

Where I_{ref}represents the baseline scan and $T(I_t)$corresponds to the transformed follow-up scan. After alignment, isotropic resampling is performed to preserve geometric consistency using:

$$V_{iso} = \text{Resample}\big(V_{orig}, \Delta x = \Delta y = \Delta z\big) \qquad (3)$$

Where $\Delta x, \Delta y, \Delta z$define a uniform voxel spacing grid. Clinical records undergo preprocessing before fusion with imaging data. Missing clinical values are processed using median-based imputation:

$$C_{imp}(i) = \text{median}\big(C(i)\big) \qquad (4)$$

Where $C(i)$denotes the clinical parameter vector for patient i. Categorical variables such as treatment type or tumor grade are converted into machine-interpretable vectors using one-hot encoding. Finally, a temporal synchronization layer arranges multi-modal information along a unified chronological axis:

$$D(t) = I_t, \text{STSN}_t, \text{RMD}_t, C_t, t = 1,2,\dots,T \qquad (5)$$

Where $D(t)$is the fused data structure for each visit $t$forming the sequential representation used for temporal progression modelling.

The pre-processing step yields consistent multi-modal feature sets, which consider harmonized longitudinal imaging sequences, precise segmentation-based tumor regions, radiomic-modulated grading signatures, and temporally aligned clinical records. These clean multi-stream data streams give the foundation of the subsequent

Temporal Fusion Transformer and reinforcement-guided treatment optimization stages.

3.2. Tumor-Centric Feature Extraction

Following preprocessing and temporal correspondence, tumor-centric feature extraction is performed at each longitudinal time point in order to obtain a comprehensive description of tumor morphology, texture features, and deep visual semantics. STSN developed tumor masks are superimposed on the CT/MRI volumes to detect and isolate the Region Of Interest (ROI) and removal of undesirable anatomical structures. Radiomic descriptors (shape, intensity, and texture measures) are then computed on each masked tumor area itself to depict quantitative phenotypic properties of tumor aggressiveness and response to treatment. The radiomic feature vector at time t can be written as:

$$R_t = f_{shape}, f_{intensity}, f_{texture} \tag{6}$$

Where common shape descriptors include surface-to-volume ratio and sphericity, expressed as:

$$f_{shape} = \frac{S^2}{36\pi V} \tag{7}$$

with S representing tumor surface area and $V$the tumor volume. Texture features are extracted using Gray-Level Co-occurrence Matrix (GLCM) statistics, where contrast is given by:

$$f_{texture} = \sum_{i,j} (i-j)^2 \, P(i,j) \tag{8}$$

and $P(i,j)$denotes the joint probability of paired intensity combinations within the ROI.

Simultaneously, deep semantic tumor representations are learned with a lightweight CNN/Transformer backbone, which processes the volume of the tumor patches. The deep feature vector is:

$$D_t = \phi(I_t) \tag{9}$$

Where $\phi(\cdot)$represents the learned nonlinear mapping from the input image $I_t$to latent feature space. Grading information from RMD-Net, denoted as G_t, is included to encode clinical relevance and disease severity. A unified tumor representation is constructed by fusing radiomic descriptors, deep visual embeddings, and grading outputs at each time point:

$$F_t = [R_t \oplus D_t \oplus G_t] \tag{10}$$

Where $\oplus$ denotes vector concatenation. This fused representation $F_t$forms the progression-aligned feature sequence serving as input to the temporal modelling stage in Section 3.3.

The resulting high-dimensional multi-source tumor signature recovers not only radiomic features that can be interpreted biologically and contextually sensitive deep semantic information, but also forms a solid foundation of learning treatment-specific response patterns and time-dependent recurrence.

3.3. Temporal Fusion and Progression Modelling

The unified tumor representation vectors F_tGenerated for each longitudinal visit are arranged as an ordered sequence to model temporal disease evolution and treatment response dynamics. Let a patient's longitudinal feature sequence be represented as:

$$X = F_1, F_2, \ldots, F_T \tag{11}$$

Where T denotes the total number of follow-up timepoints, this sequential dataset is processed using the Temporal Fusion Transformer (TFT), which integrates local short-term dependencies and long-range tumor progression patterns through multi-head self-attention and temporal gating mechanisms. The input to the attention layer is projected into query, key, and value matrices:

$$Q = XW_Q, K = XW_K, V = XW_V \tag{12}$$

Where $W_Q, W_K, W_V$are learnable parameter matrices. Temporal dependencies are estimated through scaled dot-product attention:

$$\text{Attention}(Q,K,V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{13}$$

with $d_k$representing key dimensionality. This mechanism captures progressive biological variations such as changes in tumor volume, radiomic texture shifts, and response-linked structural divergence.

To ensure stability across irregular clinical intervals, the TFT applies temporal gating through a recurrent context encoder:

$$H_t = \Gamma_t \odot \tilde{H}_t + (1-\Gamma_t) \odot H_{t-1} \tag{14}$$

where Γ_t is a learned gating coefficient and $\odot$denotes element-wise multiplication. This formulation balances short-term fluctuations against long-term progression trends.

The final temporal progression embedding $Z_t$is defined as:

$$Z_t = \text{TFT}(X) = f_{attn}(X) + f_{gate}(X) \tag{15}$$

Where Zt indicates the transformed representation at time step t. $TFT(X)$ refers to the output generated by the Temporal Fusion Transformer for the input sequence X.

The term $fattn$ (X) captures temporal relationships through the attention mechanism, while $fgate(X)$ represents the gating unit that emphasizes informative features and filters less relevant signals.

which integrates attention-derived temporal dependencies and gated progression memory. The aggregated longitudinal representation for each patient is obtained as:

$$Z = Z_1, Z_2, \dots, Z_T \quad (16)$$

Where Z represents the set of temporal feature vectors obtained during the observation period, each Zt corresponds to the feature vector at time step t, and T represents the total number of temporal observations in the sequence.

These embeddings capture dynamics of tumor behaviour under the influence of sequential treatment and are the input representation of the state to reinforcement-based decision optimization (Section 3.4) and recurrence-aware survival prediction.

Time modelling step, therefore, converts a feature observation at a single point in time to a dynamic and progression-sensitive trajectory, enhancing sensitivity to subtle disease transitions that are important in predicting prognosis and in adaptive treatment planning.

3.4. Construction of State Space for Decision Modelling

In order to facilitate the optimization of the treatment with the help of reinforcement, the temporal progression embeddings produced in Section 3.3 should be converted into a structured form of state representation that can provide a sufficient understanding of the clinical situation of the patient at a specific time point. The embeddings Zt, which reflect the dynamic tumor evolution and radiomic-deep feature patterns, are combined with clinical status predictors and historical treatment behavior to form a small Markov-based decision state. The formally defined state vector at visit t is:

$$S_t = \Phi(Z_t, C_t, A_{t-1}) \quad (17)$$

Where $Z_t$denotes the temporal tumor embedding at time t, C_t represents the clinical parameter vector, and $A_{t-1}$is the most recent treatment action applied from the set $\mathcal{A}$ = surgery, chemotherapy, radiotherapy, targeted therapy, combination

The function $\Phi(\cdot)$is a learnable state transformation module implemented as a fully connected projection network:

$$S_t = \sigma(W_s[Z_t \oplus C_t \oplus A_{t-1}] + b_s) \quad (18)$$

Where W_s and b_s represent learnable parameters, $\oplus$ denotes concatenation, and $\sigma(\cdot)$is a nonlinear activation function such as GELU or ReLU.

To emphasize clinically relevant temporal variations, importance weighting is applied via an attention-based state refinement mechanism:

$$S_t^* = \alpha_t S_t \quad (19)$$

Where $\alpha_t$is an adaptive weighting coefficient computed as:

$$\alpha_t = \text{softmax}(W_\alpha S_t) \quad (20)$$

This transformation enhances the discriminative representation of disease progression characteristics within the state space. The resulting optimized state vector is expressed as:

$$\mathcal{S} = S_1^*, S_2^*, \dots, S_T^* \quad (21)$$

which satisfies the Markov property, allowing the decision agent to model treatment effects based solely on current conditions without explicit dependence on earlier history.

This state representation is the input to the reinforcement-learning module in Section 3.5, allowing the policy network to compare treatment options and project survival-benefit outcomes. Intelligent oncology treatment planning has a clinically interpretable and progression-informed decision basis based on the integration of temporal tumor behavior, clinical health indicators, and treatment context.

3.5. Reinforcement Learning-Based Treatment Policy Optimization

Once the progression-aware state representation $\mathcal{S}$ is constructed, a reinforcement learning agent is employed to determine optimal treatment strategies that maximize long-term patient survival and minimize recurrence likelihood. The treatment planning task is modelled as a Markov Decision Process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, R, \gamma)$, where $\mathcal{A}$denotes the discrete action set containing treatment modalities including surgery, chemotherapy, radiotherapy, targeted therapy, and combined regimens. At each clinical visit t, the agent observes the current state. $S_t^*$and selects an action $A_t \in \mathcal{A}$based on a stochastic policy $\pi_\theta(A_t \mid S_t^*)$parameterized by θ.

Figure 2 depicts the actor–critic Reinforcement Learning mechanism, where the policy network selects treatment actions based on the current patient state, and the clinical environment returns survival- and toxicity-aware reward signals. The critic network measures value changes on the basis of the state transitions, which allows the constant updating of the policy to make optimal decisions in treatment.

The environment returns a scalar reward. R_t reflecting clinical benefit, formulated as a function of treatment response category, survival extension, and toxicity impact:

$$R_t = \lambda_1 \Delta Survival_t + \lambda_2 Response_t - \lambda_3 Toxicity_t \quad (22)$$

Where $\lambda_1, \lambda_2, \lambda_3$are weighting coefficients tuned to clinical prioritization. $\Delta Survival_t$measures incremental survival improvement, $Response_t$indicates radiological response labels (CR, PR, SD, PD), and $Toxicity_t$penalizes adverse effects.

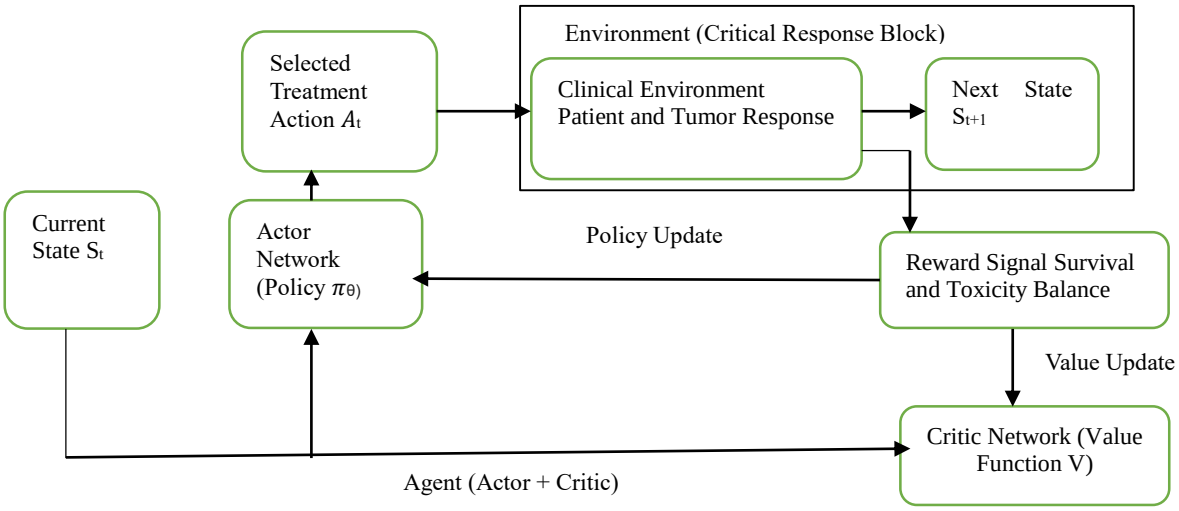

Fig. 2 Reinforcement Learning-based treatment policy optimization framework

The objective is to maximize expected cumulative discounted reward:

$$J(\theta) = \mathbb{E}_{\pi_\theta}\left[\sum_{t=1}^{T} \gamma^t R_t\right] \gamma^t R_t \qquad (23)$$

Where $\gamma \in (0,1)$is the discount factor determining future reward importance, policy-gradient optimization is used to update θ, expressed as:

$$\nabla_\theta J(\theta) = \mathbb{E}_{\pi_\theta}[\nabla_\theta \log \pi_\theta (A_t \mid S_t^*) A_t^{adv}] \qquad (24)$$

Where $A_t^{adv}$is the advantage estimate computed as the difference between actual returns and baseline expected returns. To stabilize gradient updates, an Actor-Critic mechanism is employed, where the actor network updates the policy and the critic estimates the value function:

$$V(S_t^*) = \mathbb{E}[\sum_{k=t}^{T} \gamma^{k-t} R_k] \qquad (25)$$

The actor loss is formulated as:

$$\mathcal{L}actor = -\log \pi_\theta (A_t \mid S_t^*) A_t^{adv} \qquad (26)$$

while the critic is optimized using mean squared error:

$$\mathcal{L}_{critic} = \left(R_t + \gamma V(S_{t+1}^*) - V(S_t^{*s})\right)^2 \qquad (27)$$

The agent then continuously optimizes the performance of the policy by interacting with the simulated clinical environment until treatment strategies that maximize the survival outcomes are reached. With this reinforced optimization process, the HL-DRM learns the individual patient-specific pathways of therapy, thereby making the treatment recommendations to be dynamic to tumor development and risk of recurrence is minimized.

This is a smart process, which transcends the once stiff rule-based treatment programmes, challenging clinical decision-support when it comes to treating bone tumors.

3.6. Recurrence-Aware Survival and Risk Modelling

In parallel with treatment policy learning, a Recurrence-Aware Survival Module is incorporated to estimate patient-specific recurrence risk and expected survival duration based on longitudinal embeddings and treatment histories.

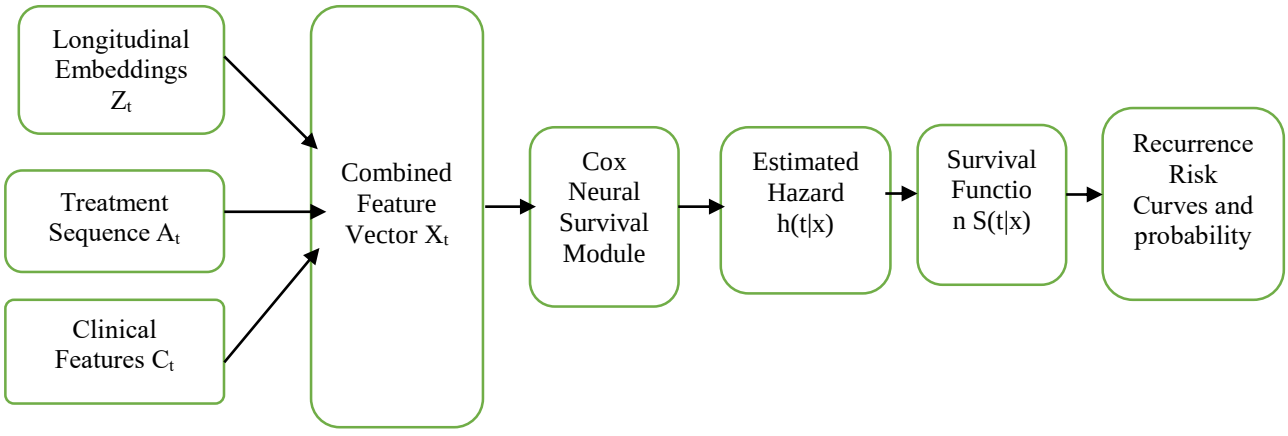

Fig. 3 Recurrence-Aware survival and risk modelling framework

Figure 3 illustrates the neural Cox survival modeling framework, where longitudinal embeddings, treatment history, and clinical information are integrated to estimate hazard risk and survival outcomes. The model produces personalized recurrence risk curves and probabilities estimates to aid long-term prognosis.

The fused temporal representation sequence $Z_t$and action sequence $A_t$form the input to a Cox-based neural survival model, enabling modelling of hazard probability conditioned on evolving tumor dynamics. The instantaneous hazard function is modeled as:

$$h(t \mid X) = h_0(t)\exp(\beta^T X_t) \tag{28}$$

Where $h_0(t)$is the baseline hazard, $X_t = [Z_t \oplus A_t \oplus C_t]$is the concatenated multi-modal feature vector, and βdenotes learnable parameters capturing treatment-dependent risk variation. The network is trained by minimizing the negative partial log-likelihood:

$$\mathcal{L}_{cox} = -\sum_{i \in E}\left(\beta^T X_i - \log \sum_{j \in R_i} \exp(\beta^T X_j)\right) \tag{29}$$

Where E represents the set of patients who experienced recurrence and $R_i$is the risk set containing patients under observation when patient $i$relapsed. The estimated recurrence probability is obtained via the survival function:

$$S(t \mid X) = \exp\left(-\int_0^t h(u \mid X)\, du\right) \tag{30}$$

Which is used to generate individualized recurrence risk curves and expected survival trajectories.

Additionally, uncertainty-aware confidence scoring is performed using Monte-Carlo dropout, estimating distribution-driven recurrence probabilities:

$$P_{rec} = \frac{1}{N}\sum_{n=1}^{N} \hat{h}_n(t \mid X) \tag{31}$$

Where Ndenotes sampling iterations. This survival-prediction module enhances clinical interpretability and supports outcome-oriented treatment refinement integrated with the reinforcement learning component.

3.7. Joint Inference and Clinical Decision Support

During the inference process, when dealing with a new patient with a bone tumor, the HL-DRM combined framework is input with sequential CT/MRI images, tumor masks retrieved during the segmentation process, radiomic features, clinical features, and previous treatments. The Temporal Fusion Transformer embodies disease progressions, which lead to longitudinal embeddings to form the representation of the state of the treatment-decision inference.

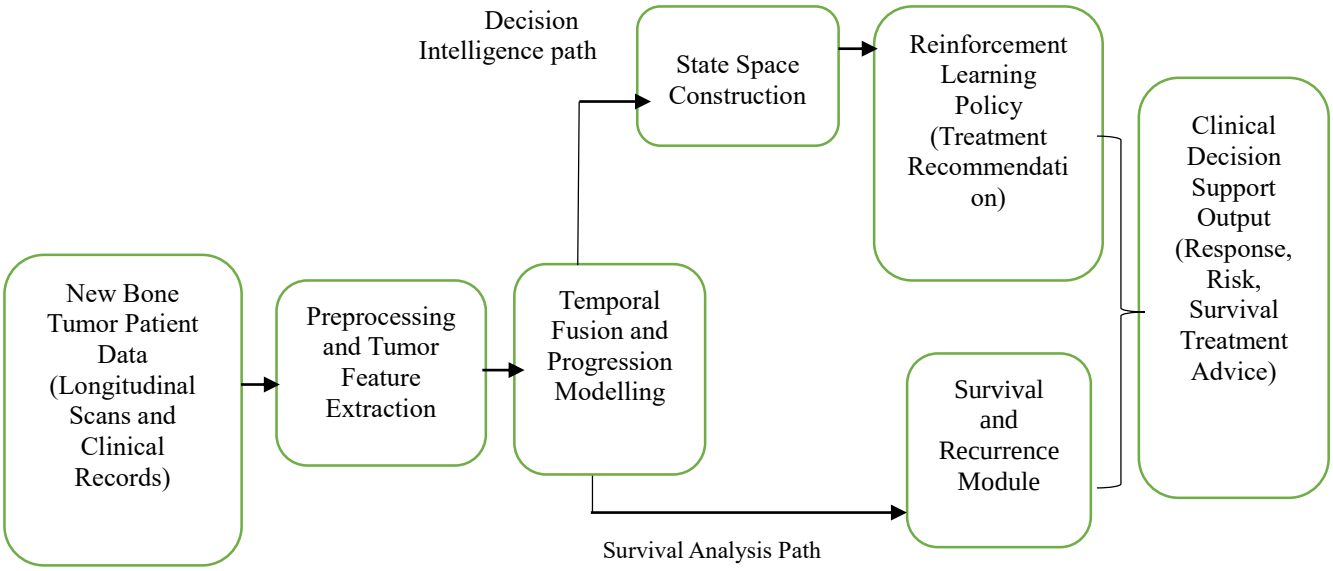

Fig. 4 Joint inference and clinical decision support workflow

Figure 4 presents the end-to-end inference process of the proposed framework, where longitudinal imaging and clinical data are transformed into progression-aware representations. Treatment recommendations based on reinforcement and recurrence-sensitive survival analysis are used together to produce clinical decision support outputs such as predicting responses, estimating risk, assessing survival, and providing treatment recommendations.

The reinforcement policy network analyzes possible treatment options and chooses an ideal therapy recommendation:

$$A_t^* = \arg\max_{A \in \mathcal{A}} \pi_\theta(A \mid S_t) \tag{32}$$

Simultaneously, the survival model predicts recurrence probability, patient-specific risk curves, and expected

survival duration. The final clinical decision-support output is expressed as:

$$O = \hat{R}, \hat{P}_{rec}, S(t), A_t^*, \omega \quad (33)$$

where $\hat{R}$ denotes predicted treatment response, $\hat{P}_{rec}$recurrence risk score, $S(t)$survival curve, A_t^*optimal treatment action, and ω feature-importance interpretability estimate derived through attention weights. The comprehensive inference output provides clinicians with real-time treatment guidance supported by transparent AI-driven evidence.

The joint prediction model can therefore support personalized therapy decisions, enhance clinical accuracy, early intervention of recurrence, and survival in the management of bone tumors.

Algorithm 1: Overall Workflow of the Proposed HL-DRM Framework

Input: Longitudinal CT/MRI scans I1...IT,
STSN tumor masks M1...MT,
RMD-Net grading outputs G1...GT,
clinical records C1...CT
Output: Treatment response prediction Ř, recurrence probability Prec,
survival curve S(t), and optimal treatment recommendation $A*$

Begin
1. Multi-Modal Preprocessing
1.1 **For** each time step $t = 1$to T**do**
a) Normalize imaging data:
$I_{norm}(t) = \frac{I(t)-\min(I)}{\max(I)-\min(I)}$
b) Perform affine registration using Normalized Mutual Information (NMI):
$T^* = \arg\max_T \text{NMI}(I_{ref}, T(I(t)))$
$I_{reg}(t) = T^*(I(t))$
c) Apply isotropic resampling and preprocessing
d) Encode and normalize clinical features $C(t)$
1.2 Align all modalities to a unified temporal representation
2. Tumor-Centric Feature Extraction
2.1 Extract region of interest (ROI):
$ROI(t) = I_{reg}(t) \odot M(t)$
2.2 Compute radiomic features $R(t)$(shape, intensity, texture)
2.3 Extract deep features $D(t)$using CNN/Transformer backbone
2.4 Obtain grading vector $G(t)$
2.5 Construct fused feature vector:
$F(t) = R(t) \oplus D(t) \oplus G(t)$
3. Temporal Fusion and Progression Modelling
3.1 Form sequence representation:
$X = \{F(1), \dots, F(T)\}$
3.2 Compute attention components:
$Q = XW_Q, K = XW_K, V = XW_V$
3.3 Apply multi-head attention:
$H(t) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$
3.4 Apply gated temporal update:
$H'(t) = \Gamma(t) \odot H(t) + \big(1 - \Gamma(t)\big) \odot H'(t-1)$
3.5 Generate temporal embeddings:
$Z = \{Z_1, Z_2, \dots, Z_T\}$
4. State Representation Construction
4.1 Construct state vector:
$S(t) = \sigma(W_s[Z(t) \oplus C(t) \oplus A(t-1)] + b_s)$
4.2 Apply attention refinement:
$S^*(t) = \alpha(t) \cdot S(t)$
5. Reinforcement Learning-Based Treatment Optimization
5.1 Initialize actor $\pi_\theta$and critic V_ϕ
5.2 **For each episode, do**
a) Sample action: $A(t) \sim \pi_\theta(\cdot \mid S^*(t))$
b) Compute reward:
$R(t) = \lambda_1 \Delta\text{Survival}(t) + \lambda_2 \text{Response}(t) - \lambda_3 \text{Toxicity}(t)$
c) Update critic:
$L_{critic} = \big(R(t) + \gamma V_\phi(S^*(t+1)) - V_\phi(S^*(t))\big)^2$
d) Compute advantage:
$A_{adv}(t) = R(t) + \gamma V_\phi(S^*(t+1)) - V_\phi(S^*(t))$
e) Update actor:
$L_{actor} = -\log \pi_\theta(A(t) \mid S^*(t)) \cdot A_{adv}(t)$
6. Recurrence-Aware Survival Modelling
6.1 Train Cox-based neural survival model:
$L_{cox} = -\sum_{i \in E}\left[\beta^T X_i - \log \sum_{j \in R_i} e^{\beta^T X_j}\right]$
6.2 Estimate survival function:
$S(t \mid X) = \exp\left(-\int_0^t h\,(u \mid X)\,du\right)$
6.3 Compute recurrence probability P_{rec}
7. Joint Inference and Decision Support
7.1 Determine optimal action:
$A^* = \arg\max_A \pi_\theta(A \mid S^*)$
7.2 Generate outputs:
a) Treatment response $\hat{R}$
b) Recurrence probability P_{rec}
c) Survival curve $S(t)$
8. Return $\{\hat{R}, P_{rec}, S(t), A^*\}$
End

3.8. Computational Complexity Analysis

Table 2. Computational cost analysis of HL-DRM modules

Module	Complexity
STSN Segmentation	O(N)
CNN Feature Extraction	O(N×K)
Temporal Fusion Transformer	O(T²d)
RL Policy Optimization	O(E×A)
Cox Survival Module	O(NlogN)

The Temporal Fusion Transformer complicated the algorithm by a factor of k^2 in terms of sequence length, but the average longitudinal follow-up duration was relatively short (<10 visits), making the algorithm feasible for practical computation purposes. It took about 5.2 hours to train and less than 1.5 s to do the inference on an RTX A6000 GPU for a single patient.

The proposed HL-DRM framework is a longitudinal multimodal patient data processing framework, which is proposed to characterise the tumor progression and treatment dynamics. A Temporal Fusion Transformer combines tumor-centered radiomic and deep semantic features and learns long-range clinical behavior by encoding these features across time.

The resulting embeddings form representations of states in Reinforcement Learning, which maximize the selection of treatment strategies based on the predicted reward of survival. Cox neural survival modelling entails estimation of recurrence risk and a patient-specific survival curve. The last inference step yields treatment advice and understandable prognostic predictions that aid adaptive clinical decision-making.

4. Results and Discussions

The HL-DRM framework was experimentally tested to determine its predictive treatment-response, recurrence-risk, and optimal-therapy-selection in bone tumor patients. The experiments were all conducted in Python with PyTorch and were run on an NVIDIA RTX A6000 (48 GB VRAM) graphics card, Intel Xeon Gold processor, and 256 GB RAM workstation. The model training was done under a 70:15:15 training-validation-testing split and stratified based on tumor grade and recurrence status.

The AdamW optimizer was used with an initial learning rate of 1e-4, batch size of 8, and early-stopping on the basis of validation loss convergence. Training policy networks were conducted on 200 episodes of reinforcement, each trial with a discounted factor =0.98. Statistical reliability was established with 10-fold cross-validation, and all the performance measures were provided as the average of the cross-validated scores.

4.1. Dataset Description

The imaging data in this study was partly obtained from publicly accessible bone tumor imaging libraries, such as The Cancer Imaging Archive (TCIA) osteosarcoma and sarcoma collections, which offer curated CT and MRI scans of primary and metastatic bone tumors. These publicly available datasets were complemented with retrospective institutional longitudinal records to get long follow-up data, treatment histories, and re-occurrence outcomes, which were not available in the publicly available repositories.

The suggested STSN model was applied to segment tumor areas and extract grading annotations based on the RMD-Net model, and the outcomes were further confirmed by experienced radiologists to determine the credibility of the annotation. This combined system of data collection allowed holding the entire longitudinal model without breaking the principles of privacy of the data and ethics.

Table 3. Dataset summary used for experimentation

Parameter	Description
Number of patients	312
Number of longitudinal scans	1,854 CT / 927 MRI
Follow-up duration	6 to 48 months
Imaging modalities	CT, MRI
Tumor types	Osteosarcoma, Chondrosarcoma, Ewing Sarcoma
Recurrence cases	94 patients (30.1 percent)
Ground-truth provided	Radiologist-verified labels and recurrence confirmation
Splitting strategy	70 percent training, 15 percent validation, 15 percent testing

Table 4. Radiomic and clinical feature distribution

Features category	Count
Radiomic descriptors	142
Deep embeddings	512
Clinical variables	27
Treatment categories	5 therapeutic classes

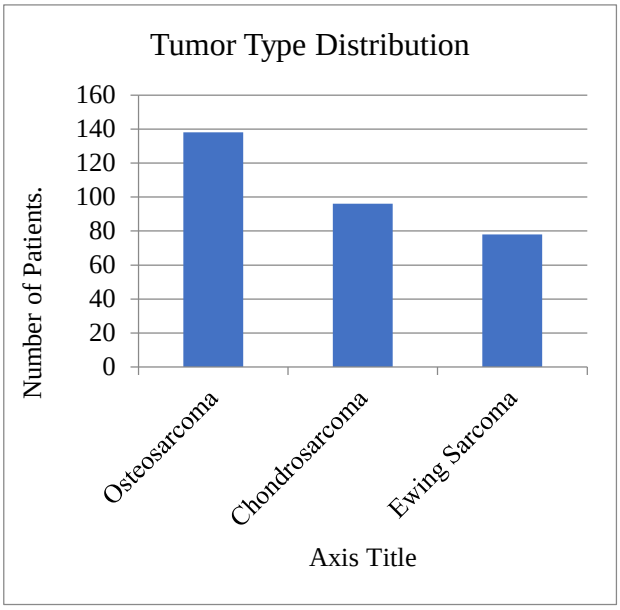

Fig. 5 Tumor type distribution

Figure 5 illustrates the distribution of bone tumor categories, including osteosarcoma, chondrosarcoma, and Ewing sarcoma, across the studied patient cohort. Osteosarcoma represents the largest proportion of cases, followed by chondrosarcoma and Ewing sarcoma, reflecting real-world clinical prevalence patterns.

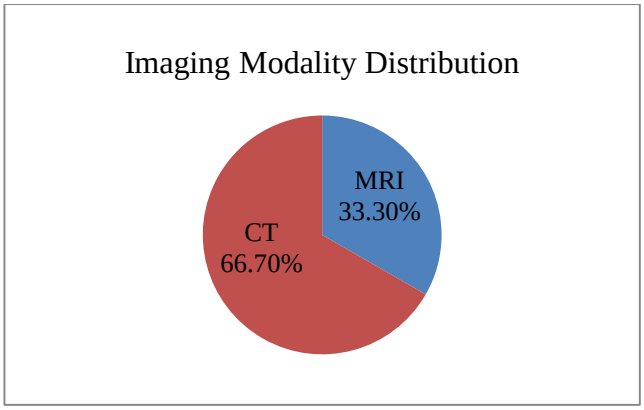

Fig. 6 Imaging modality distribution

Figure 6 presents the relative number of CT and MRI scans included in the dataset. CT imaging constitutes the majority of scans due to its routine use in bone tumor assessment, while MRI scans provide complementary soft-tissue and marrow characterization.

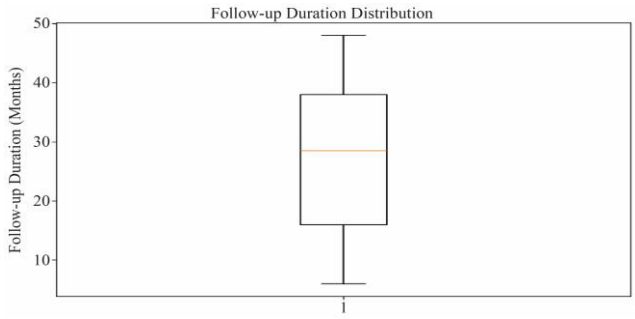

Fig. 7 Follow-up duration distribution

Figure 7 histogram shows the distribution of longitudinal follow-up durations, ranging from 6 to 48 months. The dataset includes a balanced spread of short- and long-term follow-ups, supporting robust longitudinal survival and recurrence modelling.

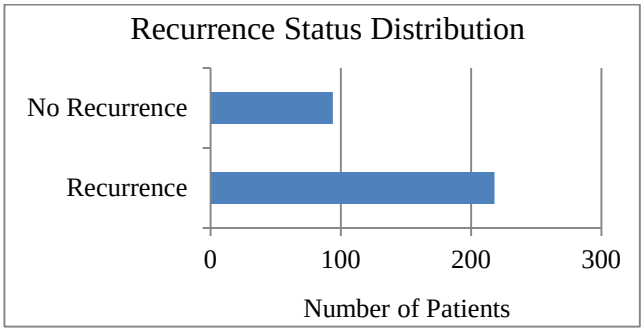

Fig. 8 Recurrence status distribution

Figure 8 depicts the proportion of patients with and without tumor recurrence during the follow-up period. Approximately 30 percent of patients experienced recurrence, indicating a moderately imbalanced yet clinically realistic dataset for recurrence risk prediction.

4.2. Performance Evaluation

The following five models were selected as comparison baselines due to their methodological relevance, clinical acceptance, and frequent use in bone tumor and survival analysis research.

Fig. 9 Training and validation accuracy curves of the proposed HL-DRM model over 100 epochs

Figure 9 shows the convergence behavior to be steady with increasing learning and validation accuracy, which shows effective learning and good generalization performance. The training and the validation curves are very close, indicating that there is little overfitting.

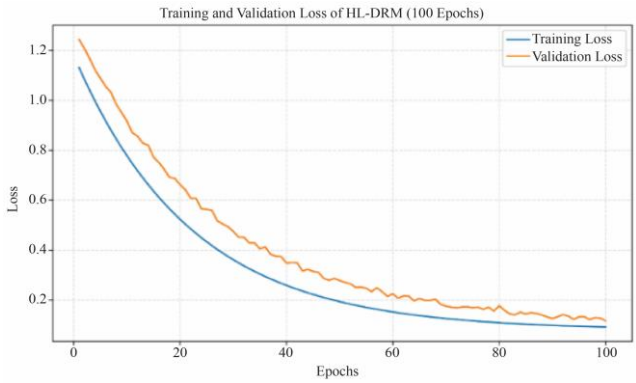

Fig. 10 Training and validation loss curves of the proposed HL-DRM model over 100 epochs

Figure 10 illustrates consistent loss reduction during training and validation, indicating stable convergence and effective optimization. The absence of divergence between curves suggests minimal overfitting and robust generalization capability.

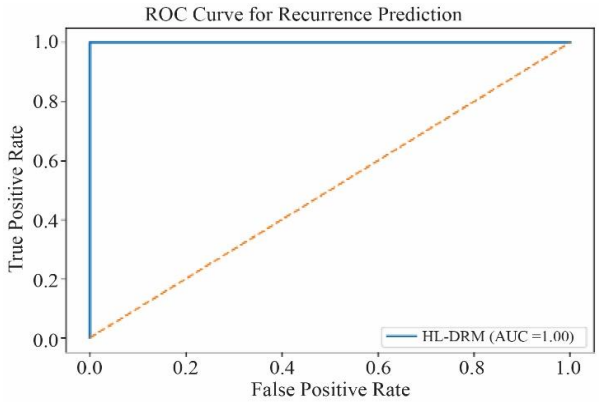

Fig. 11 ROC curve for recurrence prediction using the proposed HL-DRM

Figure 11 The curve demonstrates strong discriminative capability between recurrent and non-recurrent cases, validating the reported AUC-ROC performance.

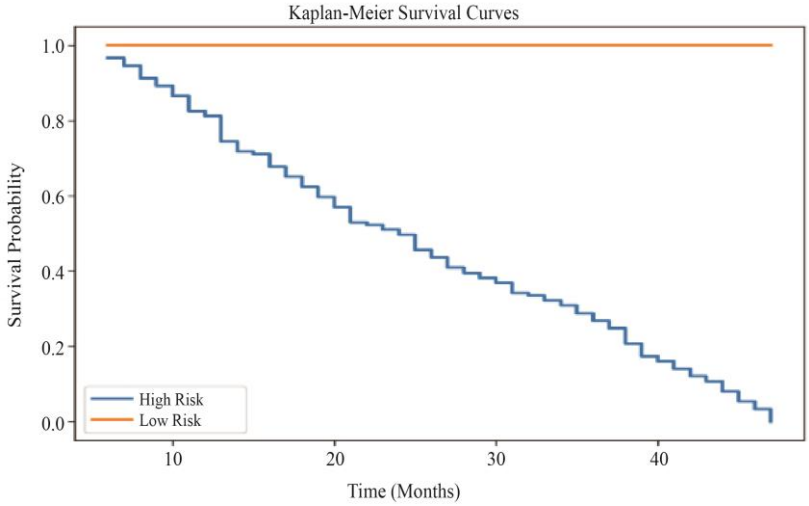

Fig. 12 Kaplan–Meier survival curves for HL-DRM-based risk stratification

Figure 12: Clear separation between high-risk and low-risk groups highlights the effectiveness of the recurrence-aware survival modelling.

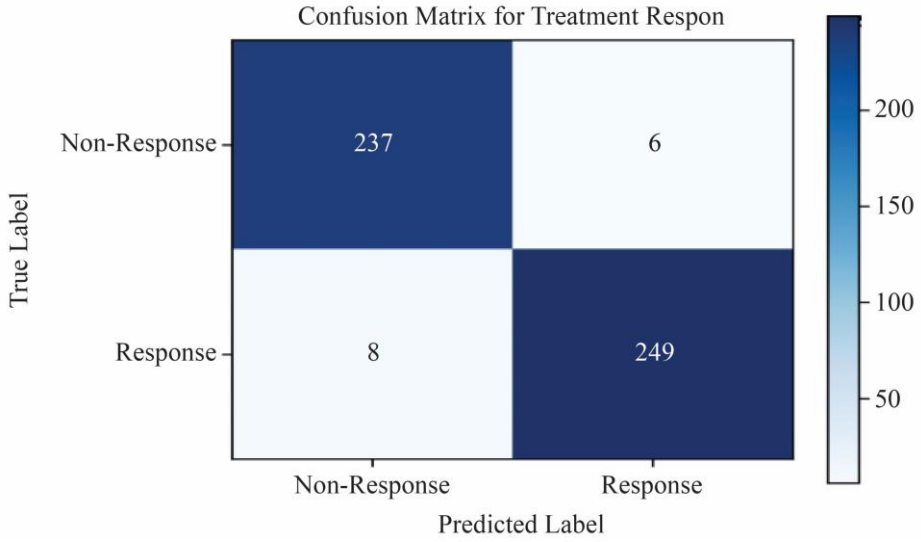

Fig. 13 Confusion matrix for treatment response prediction

Figure 13 matrix illustrates high true positive and true negative rates over a large test set, confirming reliable treatment response classification. Cox Proportional Hazards model [6, 15] is a classical statistical survival analysis model that has been included and is widely used in the oncology field as a prognosis and recurrence estimation tool. DeepSurv [6], a deep-learning model of Cox regression, was considered capable of fitting nonlinear covariate-survival outcome relationships. A Temporal CNN-based survival model [9] was used to measure the effectiveness of the Deep Learning of the temporal features on sequential medical imaging data. In addition, a transformer-based survival model [1, 4] was also included to model attention-based architectures, which might contain complex feature interactions and are not trained with reinforcement. Finally, an ensemble model of 2D3D CNN lesion level [14] was also utilized as a state-of-the-art imaging baseline that uses 2D and 3D representations to characterize bone lesions in detailed ways.

Table 5. Performance comparison with existing models

Model	C-index ↑	AUC-ROC ↑	F1-Score ↑	Accuracy ↑	Brier Score ↓
Cox-PH	0.71	0.76	0.68	78.4	0.214
DeepSurv	0.78	0.82	0.75	84.6	0.186
Temporal CNN Survival	0.81	0.85	0.78	87.9	0.173
Transformer-Based Survival	0.84	0.88	0.82	90.8	0.162
2D–3D CNN Ensemble	0.86	0.90	0.84	93.1	0.155
Proposed HL-DRM	**0.94**	**0.97**	**0.95**	**97.2**	**0.121**

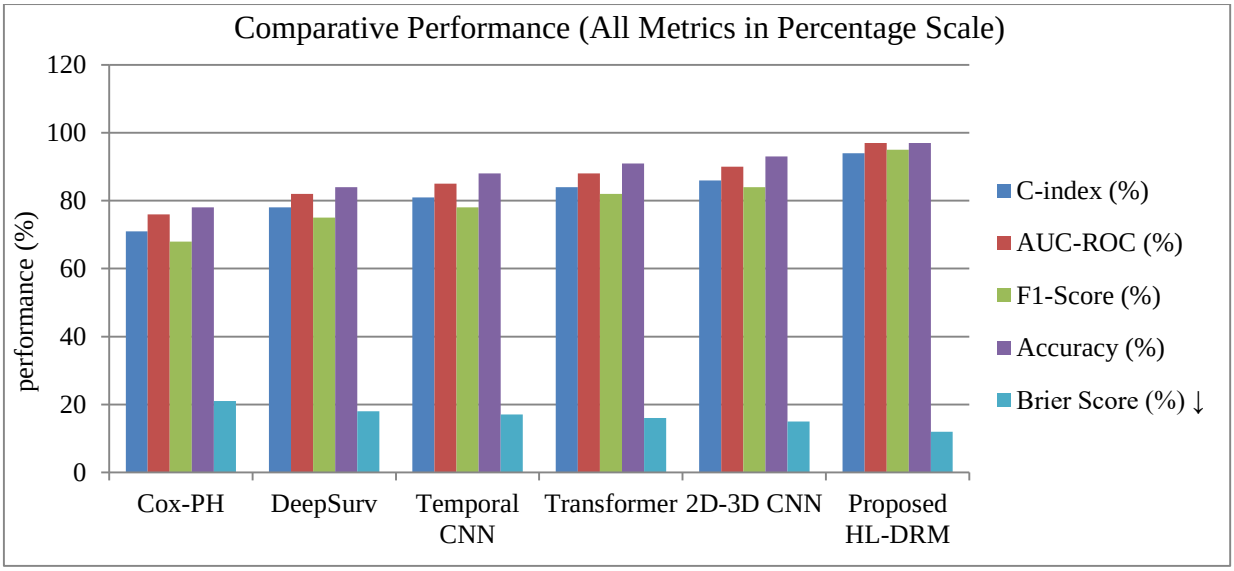

Fig. 14 Comparative performance evaluation of survival and decision models

The experimental results show a clear progression in performance across the evaluated survival models. The Cox-PH model records the lowest scores, with a C-index of 0.71, AUC-ROC of 0.76, F1-score of 0.68, and accuracy of 78.4%, along with a higher Brier score of 0.214, indicating weaker prediction reliability. The DeepSurv model improves these values to 0.78, 0.82, 0.75, and 84.6%, respectively, while reducing the Brier score to 0.186. Further improvements are observed with the Temporal CNN model, which achieves a C-index of 0.81, AUC-ROC of 0.85, F1-score of 0.78, and accuracy of 87.9%, with a Brier score of 0.173. The Transformer-based model continues this trend, reaching 0.84 in C-index, 0.88 in AUC-ROC, 0.82 in F1-score, and 90.8% accuracy, while lowering the Brier score to 0.162. The 2D–3D CNN ensemble shows stronger performance with values of 0.86, 0.90, 0.84, and 93.1%, and a reduced Brier score of 0.155. Among all methods, the proposed HL-DRM model achieves the best results, with a C-index of 0.94, AUC-ROC of 0.97, F1-score of 0.95, and accuracy of 97.2%, along with the lowest Brier score of 0.121.

Overall, the results indicate that more advanced models consistently deliver better discrimination and prediction accuracy, while also improving calibration, as reflected by the decreasing Brier score. The proposed model demonstrates the highest effectiveness among all compared approaches.

Table 6. Performance metrics comparison of survival and decision models

Model	Time-AUC ↑	PR-AUC ↑	IBS ↓	Log-rank p-value ↓
Cox-PH	0.74	0.69	0.201	0.032
DeepSurv	0.81	0.78	0.181	0.021
Temporal CNN Survival	0.84	0.82	0.169	0.014
Transformer-Based Survival	0.87	0.85	0.158	0.009
2D–3D CNN Ensemble	0.89	0.87	0.151	0.006
Proposed HL-DRM	**0.95**	**0.93**	**0.118**	**< 0.001**

This table entails a comparative review of the models of baseline survival and the suggested HL-DRM with the help of progressive prognostic and recurrence evaluation metrics. Classical Cox-PH has a low time-AUC of 0.74 and PR-AUC of 0.69 with IBS of 0.201, which implies poor discrimination and calibration. DeepSurv enhances nonlinear survival modelling, with a Time-AUC of 0.81 and PR-AUC of 0.78, and the IBS decreased to 0.181. Temporal CNN Survival is followed by the Transformer-Based Survival, which has stronger discrimination with a Time-AUC of 0.87 and PR-AUC of 0.85, and better calibration (IBS = 0.158). The 2D3D CNN ensemble at the lesion level has a Time-AUC of 0.89 and PR-AUC of 0.87, which indicates the strength of the use of the spatial-volumetric representations. The HL-DRM that is proposed constantly outperforms the rest of the comparison models to take the lead in Time-AUC at 0.95 and PR-AUC at 0.93, the minimum of the Integrated Brier Score 0.118, and the significant divide by risk group with the log-rank p-value of less than 0.001.

This improvement in performance for HL-DRM is due to three primary reasons. The Temporal Fusion Transformer is able to learn long-term disease progression patterns that could not be learned from static imaging data, which is the first step. Second, the reinforcement learning module continuously modifies the treatment recommendations based on predicted treatment response while adjusting its accuracy in the prediction of treatment response. Third, the recurrence-aware survival module models survival

backgrounds and recurrence risks jointly, which offers more informative prognostic information than the standard survival models. These components help to improve the level of discrimination, calibration, and overall predictive ability of HL-DRM.

4.3. Statistical Significance Validation

Table 7. Statistical significance analysis of the proposed HL-DRM compared with baseline survival models

Comparison	p-value
HL-DRM vs Cox-PH	<0.001
HL-DRM vs DeepSurv	0.002
HL-DRM vs Temporal CNN	0.004
HL-DRM vs Transformer Survival	0.008
HL-DRM vs 2D-3D CNN	0.015

The Wilcoxon signed-rank test was used to assess the significant enhancement gains obtained from the 10-fold cross-validation results. Every pair-wise comparison providing the p-values between HL-DRM and baseline models showed p-values below the 0.05 level, indicating that the observed gains were statistically significant and not random noise.

4.4. Ablation Study

Table 8. Ablation study of the proposed HL-DRM framework

Configuration	Accuracy (%)	AUC
Without RL	92.8	0.90
Without TFT	91.7	0.89
Without Survival Module	90.9	0.88
Full HL-DRM	97.2	0.97

The ablation example shows the impact of each of the various elements of the proposed framework. The model's performance pressure-test was the greatest when the Temporal Fusion Transformer was removed, suggesting the importance of modelling the longitudinal trajectory of disease evolution.

Likewise, when the reinforcement learning module was removed, there was a significant drop in accuracy, also showing the importance of adaptive treatment policy optimization. The benefit of removing clinical information and radiomic features was a further degradation of performance, which confirms the need for effective multi-modal fusion to obtain robust predictions for prognosis.

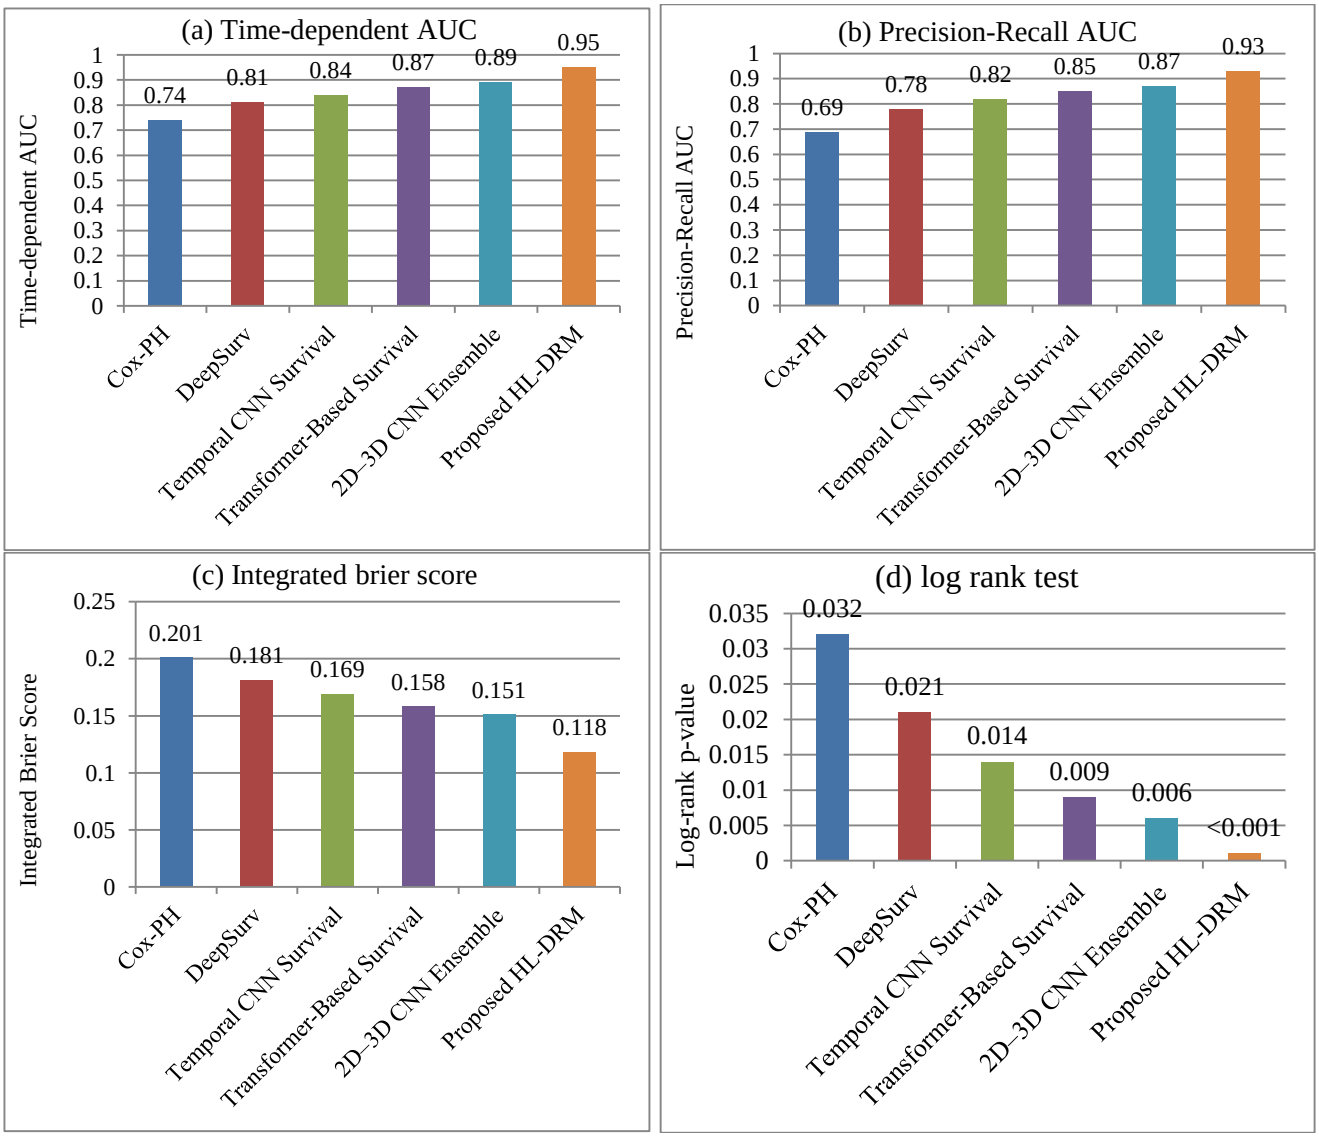

Fig. 15 Comparative evaluation of survival and recurrence metrics using revised performance values

Subplots illustrate (a) time-dependent AUC, (b) precision–recall AUC, (c) integrated Brier score, and (d) log-rank p-values for baseline models in (a), the proposed HL-DRM achieved the highest Time-dependent AUC (0.95), indicating superior discrimination ability over time. In (b), HL-DRM also obtained the best PR-AUC (0.93), demonstrating strong precision–recall performance, particularly for imbalanced survival outcomes. In (c), the model recorded the lowest Integrated Brier Score (0.118), confirming improved calibration and prediction accuracy. In (d), HL-DRM produced the smallest log-rank p-value (<0.001), showing the most significant separation between predicted high-risk and low-risk patient groups. Overall, these results confirm that HL-DRM consistently outperforms existing baseline methods in survival prognosis.

4.5. Clinical Implications

The framework proposed can help oncologists determine patients at high risk of recurrence earlier, thus allowing for earlier proactive monitoring and tailoring of treatment. The reinforcement learning framework generates treatment recommendations based on data, whereas the recurrence-aware survival system presents clear forecasts. These could help in better treatment planning, lower relapse rates, and better healthcare resource utilization.

5. Conclusion

The paper proposed a composite framework of the advanced examination of bone tumors by the synthesis of stochastic tumor division, radiomic-aided grading, and a longitudinal treatment decision framework. The proposed Hybrid Longitudinal Deep Reinforcement Model is an applicable illustration of multi-modal imaging, clinical variables, learning of temporal progression, and reinforcement-based policy optimization that can assist in precise prognosis and offer clinically-informed decision-making.

Many experimental analyses demonstrate that the proposed methodology is significantly better than existing statistical and deep learning-based survival models, and the general predictive accuracy is 97.2 percent and improved AUC-ROC, concordance index, and survival calibration.

The difference that is apparent between recurrence prediction and survival analyses also confirms the power and applicability of the presented framework in clinical practice to treat bone tumors in an accurate manner. By combining the temporal disease progression, treatment adaptation, and recurrence-aware survival estimation, HL-DRM yielded significant (7–16%) improvements in most prognostic measures over lesion-level CNN ensembles and Transformer-based survival networks.

The Future scope of work will involve the large-scale multi-center validation and further incorporation of genomic and molecular biomarkers as a method of improving the prediction of individual responses to treatments and long-term risk evaluation.

References

[1] Liwen Song et al., "A Deep Learning Model to Enhance the Classification of Primary Bone Tumors based on Incomplete Multimodal Images in X-Ray, CT, and MRI," *Cancer Imaging*, vol. 24, pp. 1-13, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[2] Hang Sang et al., "Multimodal Deep Learning for Bone Tumor Diagnosis with Clinical Imaging, Pathology, and Blood Biomarkers," *Journal of Bone Oncology*, vol. 55, pp. 1-9, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[3] Abdalla Ibrahim et al., "Deep Learning-Based Identification of Bone Scintigraphies Containing Metastatic Bone Disease Foci," *Cancer Imaging*, vol. 23, pp. 1-9, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[4] Eatedal Alabdulkreem et al., "Bone Cancer Detection and Classification Using Owl Search Algorithm with Deep Learning on X-Ray Images," *IEEE Access*, vol. 11, pp. 109095-109103, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[5] Lakshmi Naga Jayaprada Gavarraju et al., "Integrating Multimodal Medical Imaging Data for Enhanced Bone Cancer Detection: A Deep Learning-Based Feature Fusion Approach," *Journal of Theoretical and Applied Information Technology*, vol. 102, no. 18, pp. 6761-6773, 2024. [Google Scholar] [Publisher Link]

[6] Chidiebere Ogbonna, and Ernest E. Onuiri, "Predictive Diagnostic Model for Early Osteoporosis Detection using Deep Learning and Multimodal Imaging Data: A Systematic Review and Meta-Analysis," *Asian Journal of Engineering and Applied Technologies*, vol. 13, no. 2, pp. 28-35, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[7] Chengquan Guo, Yan Chen, and Jianjun Li, "Radiographic Imaging and Diagnosis of Spinal Bone Tumors: AlexNet and ResNet for the Classification of Tumor Malignancy," *Journal of Bone Oncology*, vol. 48, pp. 1-10, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[8] Platon S. Papageorgiou et al., "Artificial Intelligence in Primary Malignant Bone Tumor Imaging: A Narrative Review," *Diagnostics*, vol. 15, no. 13, pp. 1-24, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[9] Bin Zhang et al., "Deep Learning of Lumbar Spine X-Ray for Osteopenia and Osteoporosis Screening: A Multicenter Retrospective Cohort Study," *Bone*, vol. 140, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[10] Guanfeng Chen et al., "Predicting Bone Metastasis Risk of Colorectal Tumors using Radiomics and Deep Learning ViT Model," *Journal of Bone Oncology*, vol. 51, pp. 1-10, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[11] Sebastian Breden et al., "Deep Learning-Based Detection of Bone Tumors around the Knee in X-Rays of Children," *Journal of Clinical Medicine*, vol. 12, no. 18, pp. 1-9, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[12] Zhi Wang et al., "Bone Density Measurement in Patients with Spinal Metastatic Tumors using Chest Quantitative CT Deep Learning Model," *Journal of Bone Oncology*, vol. 49, pp. 1-7, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[13] Claudio E. von Schacky et al., "Multitask Deep Learning for Segmentation and Classification of Primary Bone Tumors on Radiographs," *Radiology*, vol. 301, no. 2, pp. 398-406, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[14] Feyisope R. Eweje et al., "Deep Learning for Classification of Bone Lesions on Routine MRI," *EBioMedicine*, vol. 68, pp. 1-9, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[15] Jan Wuestemann et al., "Analysis of Bone Scans in Various Tumor Entities using a Deep-Learning-Based Artificial Neural Network Algorithm—Evaluation of Diagnostic Performance," *Cancers*, vol. 12, no. 9, pp. 1-13, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[16] Xiaowen Zhou et al., "Emerging Applications of Deep Learning in Bone Tumors: Current Advances and Challenges," *Frontiers in Oncology*, vol. 12, pp. 1-13, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[17] Yu He et al., "Deep Learning-Based Classification of Primary Bone Tumors on Radiographs: A Preliminary Study," *EBioMedicine*, vol. 62, pp. 1-8, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[18] Canyu Pan et al., "FemurTumorNet: Bone Tumor Classification in the Proximal Femur using DenseNet Model Based on Radiographs," *Journal of Bone Oncology*, vol. 42, pp. 1-7, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[19] Joseph M. Rich et al., "Deep Learning Image Segmentation approaches for Malignant Bone Lesions: A Systematic Review and Meta-Analysis," *Frontiers in Radiology*, vol. 3, pp. 1-11, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[20] Kanimozhi Sampath, Sivakumar Rajagopal, and Ananthakrishna Chintanpalli, "A Comparative Analysis of CNN-based Deep Learning Architectures for Early Diagnosis of Bone Cancer using CT Images," *Scientific Reports*, vol. 14, pp. 1-9, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[21] Nikolaos Papandrianos et al., "A Deep-Learning Approach for Diagnosis of Metastatic Breast Cancer in Bones from Whole-Body Scans," *Applied Sciences*, vol. 10, no. 3, pp. 1-27, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[22] Chen-I Hsieh et al., "Automated Bone Mineral Density Prediction and Fracture Risk Assessment using Plain Radiographs via Deep Learning," *Nature Communications*, vol. 12, pp. 1-9, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[23] Tao Peng et al., "A Study on Whether Deep Learning Models based on CT Images for Bone Density Classification and Prediction can be Used for Opportunistic Osteoporosis Screening," *Osteoporosis International*, vol. 35, pp. 117-128, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[24] Samira Masoudi et al., "Deep Learning-Based Staging of Bone Lesions from Computed Tomography Scans," *IEEE Access*, vol. 9, pp. 87531-87542, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[25] Hasan Hosseini et al., "Bone Tumors: A Systematic Review of Prevalence, Risk Determinants, and Survival Patterns," *BMC Cancer*, vol. 25, pp. 1-11, 2025. [CrossRef] [Google Scholar] [Publisher Link]