

CLARITY: Medical World Model for Guiding Treatment Decisions by Simulating Context-Aware Disease Trajectories

Tianxingjian Ding, Yuanhao Zou, Chen Chen, Mubarak Shah, and Yu Tian

Institute of Artificial Intelligence, University of Central Florida
<https://dingtianxingjian.github.io/clarity-project-page/>

Abstract. Clinical decision-making in oncology requires forecasting how disease evolves under treatment, yet most AI systems remain static predictors that cannot model longitudinal, treatment-conditioned progression. Although generative and world models have demonstrated strong capabilities in general domains, their adaptation to medicine remains limited and insufficient for capturing complex, treatment-induced physiological dynamics across temporal scales. To address these gaps, we introduce CLARITY, a medical world model that enables counterfactual simulation of treatment-conditioned disease trajectories for clinical decision-making. By jointly encoding imaging-derived latent states, temporal intervals that capture irregular follow-ups, and patient-specific clinical context, CLARITY learns smooth and interpretable representations of disease progression, allowing the model to simulate how alternative treatments reshape future disease dynamics. Because treatment optimization is inherently sequential and uncertain, requiring evaluation of long-term outcomes across multiple possible interventions, we further propose an entropy-regularized, computationally efficient long-horizon prediction-to-decision framework that plans treatment strategies over imagined disease trajectories and iteratively refines therapy proposals through survival-aware feedback, forming a closed-loop simulation-to-decision framework for treatment planning. CLARITY achieves state-of-the-art performance in treatment planning and survival prediction across three cancer datasets, including two brain tumor cohorts (MUGlioma-Post and zero-shot on UCSF-ALPTDG) and one breast cancer dataset (ISPY-2), demonstrating strong generalization across cancer types while consistently outperforming prior generative methods and medical-domain large language model baselines.

Keywords: Medical world model · Medical imaging · Clinical decision making

1 Introduction

Clinical decision-making in oncology requires forecasting how a patient’s disease evolves under treatment over time, a process characterized by substantial




Fig. 1: Conceptual Overview of CLARITY. Conventional clinical AI systems (*left*) map imaging observations to a single predicted outcome, limiting their ability to evaluate alternative treatment strategies. In contrast, CLARITY (*right*) adopts a simulation-to-decision framework: a latent world model (Actor) simulates multiple “what-if” disease trajectories conditioned on clinical context, temporal intervals, and candidate therapies. These simulated trajectories are evaluated through survival-aware feedback, and the resulting signals iteratively refine therapy proposals (orange arrow), enabling closed-loop optimization of treatment strategies.

uncertainty. While modern AI systems achieve strong performance in static outcome prediction [10, 11, 20, 28–30, 36, 38, 39, 43, 51, 52], they remain fundamentally limited in modeling treatment-conditioned disease trajectories across longitudinal follow-ups. However, real-world clinical decisions depend not only on predicting outcomes, but on anticipating how specific interventions reshape future physiological states. This makes treatment planning fundamentally a sequential decision-making problem, where clinicians must reason over potential disease trajectories under alternative interventions (e.g., continuing temozolomide therapy, introducing targeted agents, or escalating to salvage treatments), as illustrated in Fig. 1. Selecting an optimal treatment strategy is, therefore, an inherently difficult counterfactual task: clinicians must forecast how disease may evolve under therapies that were never observed; yet such prognostication remains highly uncertain and challenging, even for human experts [17, 34]. Foundation models such as large language models (LLMs), despite impressive generalization capabilities [1, 21, 41], are not designed for structured, temporally grounded disease forecasting. Dynamic clinical reasoning requires models that (i) represent physiological state transitions, (ii) incorporate heterogeneous patient-specific contexts, and (iii) support interpretable decision-making processes. Importantly, effective decision support requires evaluating counterfactual treatment strategies—predicting how the disease would evolve under interventions that have not yet been applied.

World Models (WMs) offer a promising paradigm for such tasks. By learning structured latent dynamics, WMs enable forward simulation of future states and support planning over imagined trajectories [3, 4, 7, 18]. While highly successful in robotics and control, their application to medicine remains *limited*. Clinical data introduce unique challenges, including irregular time intervals, high inter-patient heterogeneity, multimodal conditioning, and stringent interpretability

requirements. Current medical generative approaches often forecast future clinical states by synthesizing raw images. However, effective clinical decision-making relies less on surface-level appearance generation and more on extracting robust visual features that accurately capture longitudinal disease progression. Because most current systems treat visual feature extraction and treatment optimization as completely disjointed stages, they lack an explicit mechanism to connect sequential visual state transitions with specific therapeutic interventions. Consequently, these models cannot support counterfactual reasoning about how different treatments may alter future disease dynamics. We therefore introduce CLARITY, a medical world model that directly links long-term visual forecasting with optimal treatment planning (Fig. 1). Unlike traditional predictive models that produce a single outcome estimate, CLARITY simulates multiple potential disease trajectories conditioned on candidate therapies and contextual factors. By modeling disease evolution directly within a latent space rather than synthesizing images, the framework captures continuous physiological transitions while enabling efficient long-horizon reasoning and counterfactual evaluation of treatment strategies.

A central design principle of CLARITY is explicit conditioning on **temporal and clinical contexts**. Continuous time intervals are encoded as embeddings to distinguish short-term response from long-term evolution, while multimodal patient attributes, including genomics, demographics, and therapeutic history are integrated to enable personalized trajectory modeling. This design allows the model to reason over irregular follow-up intervals and heterogeneous patient contexts commonly encountered in real-world clinical data.

Crucially, CLARITY bridges prediction and decision-making through a novel **Inverse Survival Evaluation**. Predicted latent rollouts are optimized under an entropy-regularized long-horizon objective and fed back into a therapy policy module, enabling iterative reassessment and refinement of treatment strategies. Instead of selecting therapies greedily from a single prediction, CLARITY performs planning over simulated disease trajectories, evaluating survival implications across candidate interventions before committing to an action. This establishes a closed-loop prediction-to-decision framework that mirrors clinical reasoning: simulate potential trajectories, evaluate survival implications, and iteratively refine treatment policies.

Overall, our contributions are threefold:

- **Treatment-conditioned disease modeling.** We model longitudinal progression as continuous physiological transitions within a visual representation space, prioritizing biological consistency and enabling counterfactual modeling of treatment-induced disease dynamics.
- **Temporal and clinical conditioning.** We encode continuous time and multimodal patient context (e.g., genomics, demographics, prior therapies) to generate longitudinally coherent and individualized predictions.
- **Prediction-to-decision integration.** We introduce an inverse survival evaluation framework that transforms latent rollouts into adaptive and interpretable treatment recommendations.

2 Related Works

World Models. The development of World Models has evolved along several paradigms. Generative-interactive models such as Genie [7] simulate controllable environments from video, while latent planning frameworks such as DreamerV3 [18] learn compact recurrent state-space models (RSSM [13]) for long-horizon control. More recently, non-reconstructive approaches such as V-JEPA [6] and MuDreamer [8] predict representations directly in latent space without pixel reconstruction. In medicine, however, world-model approaches remain limited. Existing methods typically rely on pixel-space generative modeling to predict future clinical states [24, 25, 47, 49, 50]. For example, MeWM [47] uses diffusion models to synthesize post-treatment tumor images. While these models produce visually plausible predictions, they focus on image synthesis rather than modeling treatment-conditioned disease dynamics for decision-making. Our approach differs in three key aspects. First, existing methods typically predict a single future state, while our framework performs long-horizon trajectory simulation with time-conditioned latent dynamics. Second, diffusion-based image synthesis models do not naturally incorporate heterogeneous multimodal clinical context, whereas CLARITY explicitly conditions latent transitions on patient-specific clinical information for personalized forecasting and treatment planning. Third, prior medical generative models operate in pixel space and require reconstruction, whereas CLARITY models disease evolution directly through latent state transitions.

Survival Analysis. Traditional survival analysis relies on statistical models such as Cox regression [12] and Random Survival Forests [22], which provide interpretable hazard estimation but assume proportional hazards and linear relations. Deep-learning extensions like DeepSurv [23], DeepHit [27], and Deep Survival Machines [33] relax these assumptions, learning nonlinear risk mappings from high-dimensional clinical data. Recent multimodal frameworks [16, 31] further integrate imaging and omics features for personalized prognosis. Despite these advances, most survival models remain static predictors that estimate risk from observed patient states without modeling how disease trajectories evolve under different treatment decisions. Dynamic approaches such as Latent ODE [32] and Dynamic DeepHit [26] introduce temporal continuity in survival modeling, yet they do not explicitly simulate treatment-conditioned disease trajectories or support long-horizon reasoning over alternative treatment strategies. In contrast, CLARITY integrates survival prediction with a latent world model that simulates treatment-conditioned disease evolution. By evaluating simulated trajectories under different therapies, the framework enables survival-aware planning over imagined futures rather than static outcome prediction.

3 Method

3.1 Overview

CLARITY consists of three main components: a parameter-efficient visual backbone (MRI Encoder), a Therapy Policy Agent implemented with a multimodal




Fig. 2: CLARITY’s Inference Pipeline for Direct Survival Evaluation: A frozen MRI Encoder processes the pre-imaging observation to extract a pre-treatment latent representation. In parallel, the Therapy Policies Agent (*e.g.*, GPT-5) takes the patient’s clinical context to propose multiple candidate drug combos. The Actor module (Diseases Evolution Model) then sequentially evaluates each combo one-by-one, integrating the pre-treatment latent, clinical context, temporal context, and the specific drug combo to predict a final risk score for survival analysis.

large language model (MLLM), and an Actor module (the Diseases Evolution Model). The pipeline first extracts a pre-treatment visual state z_{pre} from the input MRI using a pre-trained vision backbone adapted via LoRA, which serves as the initial state for disease trajectory simulation. Conceptually, the Actor functions as a latent world model that simulates treatment-conditioned disease state transitions.

CLARITY operates through two complementary inference modes. In **Direct Survival Evaluation** (Fig. 2), the Therapy Policy Agent proposes candidate therapies based on the patient’s clinical context, and the Actor evaluates each option by integrating z_{pre} , clinical context, temporal context, and therapy embeddings to predict a survival risk score. In **Inverse Survival Evaluation** (Fig. 4), these scores are fed back to the agent to iteratively refine therapy proposals, enabling progressive search for treatments that minimize predicted risk. During training, the Actor and LoRA adapters are optimized using paired pre- and post-treatment MRIs together with survival supervision to learn treatment-conditioned latent transitions. During inference, the MLLM is used only to generate candidate therapies, while the Actor evaluates their projected outcomes.

3.2 Therapy Policies Agent

In CLARITY, a multimodal large language model (MLLM) (*e.g.*, GPT-5 [37]) serves as the *Therapy Policy Agent* π_{MLLM} . This design allows the system to generate clinically valid treatment combinations while maintaining flexibility in exploring alternative therapy strategies. Conditioned on a high-level goal g (*e.g.*, “minimize the predicted risk score”), the agent takes as input the pre-treatment MRI x_0 and clinical context c_p , and produces structured therapy actions that satisfy medical safety constraints Ω (*e.g.*, avoiding incompatible regimens such as co-administration of *Bevacizumab* and *Temozolomide*). Formally, the MLLM

generates a candidate action set $\mathcal{A}^{(0)} = \{a_j^{(0)}\}_{j=1}^{M_0}$, where each $a_j^{(0)}$ represents a feasible treatment configuration. Each action a contains various intervention components:

$$a = \{a_{\text{chemo}}, a_{\text{radio}}, a_{\text{brachy}}, a_{\text{immuno}}, a_{\text{add}}\}, \quad (1)$$

corresponding respectively to *chemotherapy* (e.g., Temozolomide), *external radiotherapy*, *brachytherapy*, *immunotherapy*, and *additional* supportive strategies. The generation process follows guideline-informed prompting templates to ensure clinical validity and parameter consistency (drug type, dose, and schedule). Each generated therapy description is then encoded by a pretrained text encoder (e.g., [35]), yielding a dense embedding aligned with the Actor’s latent space:

$$h^{\text{drug}} = \text{Pool}(\Phi_{\text{text}}(a)) \in \mathbb{R}^d, \quad (2)$$

where d denotes the embedding dimension of the latent space.

3.3 Visual State Encoding

Let $x_0 \in \mathbb{R}^{H \times W \times D \times C}$ denote the pre-treatment 3D MRI scan. We employ a pretrained 3D vision foundation model $\mathcal{F}_\theta(\cdot)$ (e.g., [14]) to encode anatomical and pathological features from the input volume. To align the representation space with treatment-conditioned disease dynamics, we apply Low-Rank Adaptation (LoRA) [19] to the frozen backbone by injecting trainable low-rank matrices into its attention layers. This parameter-efficient adaptation allows the model to specialize to longitudinal disease modeling while preserving the general visual representations of the foundation model. The adapted encoder maps the raw imaging observation x_0 to a latent visual state:

$$z_{\text{pre}} = \mathcal{F}_{\theta+\Delta\theta}(x_0) \in \mathbb{R}^{d_{\text{pre}}}, \quad (3)$$

where $\Delta\theta$ denotes the LoRA parameters and d_{pre} is the latent representation dimension. The resulting representation z_{pre} serves as the initial disease state for the Actor module. During training, paired pre- and post-treatment MRIs allow the model to learn treatment-conditioned transitions between latent disease states, which are detailed in Sec. 3.5.

3.4 Clinical and Temporal Contexts

The Actor receives not only the visual state z_{pre} and therapy embeddings h_{drug} , but also **clinical and temporal conditioning signals** that modulate disease evolution dynamics.

Clinical Context. For each patient p , we define a structured clinical profile $\mathcal{C}_p$ containing demographic attributes (e.g., age, sex), molecular biomarkers (e.g., *IDH1/2*, *ATRX*, *1p19q* co-deletion, *MGMT* methylation), and treatment-related indicators. To obtain a unified representation of heterogeneous clinical variables, we serialize $\mathcal{C}_p$ into a textual prompt and encode it using a medical text encoder

[35]. The resulting embedding is then projected through a lightweight MLP to obtain the clinical representation

$$h^{\text{clin}} = \text{MLP}_{\text{clin}}(\Phi_{\text{text}}(\mathcal{C}_p)) \in \mathbb{R}^{d_c},$$

where d_c denotes the clinical embedding dimension. This representation conditions the latent disease dynamics, enabling the Actor to generate patient-specific trajectory predictions and survival estimates.

Temporal Context. Disease progression occurs over irregular clinical follow-up intervals. Given observations acquired at times t_{pre} and t_{post} , we encode the time gap $\Delta t = t_{\text{post}} - t_{\text{pre}}$ using sinusoidal embeddings:

$$\gamma(\Delta t) = [\sin(\omega_i \Delta t), \cos(\omega_i \Delta t)]_{i=1}^{d_t/2}, \quad (4)$$

where $\omega_i = 1/10000^{2i/d_t}$ and d_t denotes the embedding dimension. This continuous temporal encoding allows the Actor to model treatment-conditioned disease transitions across arbitrary follow-up intervals, providing explicit awareness of the elapsed time between observations.

3.5 Diseases Evolution

As shown in Fig. 2, we apply a Diseases Evolution Model as the Actor module, which models how a patient’s disease state evolves under therapy and how this evolution relates to survival outcome. It consists of two submodules: a **Post-Treatment Latent Predictor**, which forecasts post-treatment latent representations, and a **Survival Predictor**, which estimates risk and survival probabilities from latent features. The detailed architectures of Post-Treatment Latent Predictor and Survival Predictor




Fig. 3: Training Pipeline of the Actor. The **Post-Treatment Latent Predictor** consists of a N -layer self-attention Transformer to forecast post-treatment latents. The **Survival Predictor** uses a M -layer two-way cross-attention Transformer to estimate risk score and survival rate.

are shown in Fig. 3. Conceptually, the Actor functions as a world model that predicts treatment-conditioned disease state transitions.

Information Integration. To model treatment-conditioned state transitions, the Actor integrates four inputs into the Diseases Evolution Model. Specifically, the model receives pre-treatment latent z_{pre} , clinical context h^{clin} , temporal context $\gamma(\Delta t)$ and the drug combo h^{drug} . We concatenate these information inputs together as integrated representation, $[z_{\text{pre}}, h^{\text{clin}}, \gamma(\Delta t), h^{\text{drug}}]$.

Post-Treatment Latent Predictor. The post-treatment latent predictor is optimized with a latent consistency loss and a soft contrastive regularization:

$$\mathcal{L}_{\text{pred}} = \lambda_1 \mathcal{L}_{\text{latent}} + \mathcal{L}_{\text{con}}. \quad (5)$$

The state consistency term $\mathcal{L}_{\text{latent}}$ enforces consistency between the predicted latent state $\hat{z}_{\text{post}}$ and the ground-truth state z_{post} using an ℓ_1 loss:

$$\mathcal{L}_{\text{latent}} = \|\hat{z}_{\text{post}} - z_{\text{post}}\|_1 \quad (6)$$

where d denotes the dimension of the latent space. To structure the latent space, we adopt a soft-label contrastive objective $\mathcal{L}_{\text{con}}$. Rather than using binary similarity supervision, we compute pairwise treatment similarities p_{ij} from the cosine similarities of treatment text embeddings and use them as soft targets. The predicted latent similarities q_{ij} are then aligned to these targets through a temperature-scaled symmetric contrastive loss:

$$\mathcal{L}_{\text{con}} = - \sum_{i,j}^B (p_{ij} \log q_{ij} + q_{ij} \log p_{ij}), \quad (7)$$

where B is the batch size. This formulation encourages semantically similar treatments to induce nearby latent transitions, ensuring the latent space preserves the nuanced relationships of the clinical treatment space (e.g., therapies sharing similar mechanisms, such as alkylating chemotherapy agents, produce more similar disease transitions than fundamentally different interventions like radiation therapy).

Survival Predictor. To connect latent disease evolution with clinical outcomes, we introduce a Survival Predictor that jointly processes the pre-treatment latent z_{pre} and the predicted post-treatment latent $\hat{z}_{\text{post}}$. As illustrated in Fig. 3, we employ a novel *bidirectional cross-attention module* to model interactions between disease state and treatment-induced transitions, enabling the model to capture how pre-treatment characteristics influence subsequent progression. The aggregated representation produces two outputs: a one-year survival probability for calibrated outcome estimation and a continuous risk score for survival ranking. Following standard clinical survival modeling practice, we supervise these outputs using a Brier loss for calibration and a Cox partial likelihood loss for risk ordering [23]. Our ablation experiments demonstrate that combining both losses yields more reliable survival prediction compared to single-loss variants.

3.6 Inverse Survival Evaluation

Entropy-Regularized Long-Horizon Therapy Planning. Clinical treatment optimization is inherently sequential: therapy decisions influence future disease states, which in turn affect downstream survival outcomes. To capture this dependency, CLARITY formulates treatment selection as a long-horizon planning problem over imagined latent disease trajectories generated by the




Fig. 4: CLARITY’s Inverse Survival Evaluation: This diagram illustrates the iterative prediction-to-decision feedback loop. Initial risk scores from Direct Survival Evaluation (Fig. 2) are fed into the Therapy Policies Agent. The Agent then proposes updated drug combos, which the Actor Scores to generate risk estimates as the accumulated survival feedback. This process repeats, refining the therapy proposals, and after K iterations, the policy with the Lowest Risk Score is selected as the Final Action.

Actor. Rather than selecting a single intervention greedily, the framework evaluates entire treatment schedules while maintaining exploration to account for model uncertainty and clinical variability. Formally, we model therapy planning as entropy-regularized trajectory optimization. Let $\mathbf{a}_{1:H}$ denote a treatment schedule over H decision steps. The cumulative survival risk associated with the trajectory is defined as:

$$J(\mathbf{a}_{1:H}) = \sum_{t=1}^H \beta^{t-1} \hat{r}_t, \quad (8)$$

where $\hat{r}_t$ is the predicted survival risk at step t and β is a temporal discount factor. Instead of greedily minimizing J , we optimize a trajectory distribution q under a maximum-entropy objective:

$$\max_q -\mathbb{E}_q[J(\mathbf{a}_{1:H})] + \tau \mathcal{H}(q), \quad (9)$$

where $\mathcal{H}(q)$ denotes entropy and τ controls the exploration–exploitation trade-off. In CLARITY, the trajectory distribution q is implicitly parameterized by the Therapy Policy Agent implemented with an MLLM. At iteration k , the agent proposes a set of candidate therapy trajectories sampled from its conditional policy. Each trajectory is evaluated through long-horizon latent rollouts generated by the Actor, producing survival risk estimates. These risk signals are incorporated into the next policy prompt as structured feedback, shifting probability mass toward lower-risk trajectories.

This loop effectively performs approximate entropy-regularized policy improvement: survival feedback reweights candidate trajectories while the stochastic generation capability of the language model maintains diversity in the search space. Over iterations, the distribution q progressively concentrates around survival-optimal treatment strategies without prematurely collapsing exploration.

3.7 Arbitrary Policy-Conditioned Projection

CLARITY supports counterfactual “what-if” simulations by modeling dynamic treatment regimens as sequences of therapy actions $\{a_0, \dots, a_{K-1}\}$ executed at

arbitrary discrete time points $\{t_0, \dots, t_K\}$. The Actor evaluates these policies by autoregressively generating a sequence of visual states.

Initializing the trajectory with the baseline visual state ($\hat{z}_0 = z_{\text{pre}}$), at each step $k \in \{0, \dots, K-1\}$, the model applies the learned $z_{\text{pre}} \rightarrow z_{\text{post}}$ transition dynamics. Specifically, it integrates the current state $\hat{z}_k$ (acting as the pre-treatment anchor), clinical context h^{clin} , time-gap $\phi(t_{k+1} - t_k)$, and therapy embedding h_k^{drug} to predict the subsequent state $\hat{z}_{k+1}$ (representing the post-treatment outcome). The Survival Predictor then maps this updated representation to stage-specific survival estimates $[\hat{p}_{1y}^{(k+1)}, \hat{r}^{(k+1)}]$.

This recursive rollout allows CLARITY to simulate disease trajectories under arbitrary policies *without requiring paired imaging at intermediate steps*, supporting the inverse survival optimization process (qualitative examples in Sec. 4.4).

4 Experiment

4.1 Implementations and Datasets

For training and evaluation, we leverage two longitudinal brain tumor cohorts that provide both temporal imaging trajectories and treatment information, which are essential for modeling therapy-conditioned disease evolution in a world-model framework. The **MU-Glioma-Post** [48] cohort contains 203 patients and 654 MRI follow-ups with rich treatment logs and genomic annotations. We split the dataset at the *patient level* to prevent temporal leakage, using 75% of patients for training, 15% for validation, and 15% for testing. To ensure robustness, we further perform *5-fold cross-validation* over the training-validation splits and report the mean performance across folds. To evaluate generalization, we perform **zero-shot external validation** on the **UCSF-ALPTDG** [15] brain tumor dataset, which contains 298 patients with longitudinal MRI scans and survival outcomes. The model is trained solely on MU-Glioma-Post and evaluated on UCSF-ALPTDG without fine-tuning.

Additionally, we evaluate CLARITY on the ISPY-2 breast cancer dataset [44]. This cohort contains 985 patients with longitudinal MRI scans across multiple treatment timepoints, alongside extensive molecular subtyping and neoadjuvant chemotherapy records. To demonstrate extensibility beyond neuro-oncology, the model is independently trained on this cohort to capture longitudinal therapeutic responses under neoadjuvant chemotherapy.

4.2 Results on Treatment Exploration

As shown in Tab. 1, our CLARITY framework achieves state-of-the-art results, significantly outperforming all baselines across the evaluated datasets. On the MU-Glioma-Post benchmark, our approach achieves an F1-score of 57.1%, a substantial 9.2% absolute improvement over the strongest medical-specific baseline, Huatuo-Vision (46.4%). The trend continues on UCSF-ALPTDG, where our model achieves 48.9% F1, again outperforming the second-best method (44.1%). These results highlight a key finding: while general-purpose models

Table 1: Quantitative comparison on the MU-Glioma-Post, UCSF-ALPTDG (zero-shot), and ISPY-2 datasets. MU-Glioma-Post and UCSF-ALPTDG are brain tumor datasets, while ISPY-2 is a breast cancer dataset. The **best** and second-best results are highlighted in boldface and underline, respectively. * denotes re-implemented methods for fair comparison.

Method	MU-Glioma-Post (%)				UCSF-ALPTDG (%)				ISPY-2 (%)			
	Prec.	Rec.	F1	Jac.	Prec.	Rec.	F1	Jac.	Prec.	Rec.	F1	Jac.
<i>General LLMs</i>												
GPT-5 [21]	49.3	<u>50.7</u>	<u>50.0</u>	<u>33.3</u>	38.1	47.5	42.3	26.8	37.7	46.2	41.5	26.2
Claude-4.5-Sonnet [2]	48.6	38.0	41.6	26.3	<u>45.3</u>	38.6	41.7	26.3	42.8	36.5	39.4	24.5
Qwen3-VL [5, 42]	36.7	39.4	38.0	23.5	33.7	42.9	35.8	21.8	32.5	40.1	35.9	21.9
<i>Medical models</i>												
MedGPT [46]	41.6	42.1	41.9	26.5	36.7	46.3	40.9	25.7	38.2	43.5	40.7	25.6
Huatuo-Vision [9]	<u>52.3</u>	46.8	46.4	30.2	42.1	51.5	<u>44.1</u>	<u>28.3</u>	<u>48.4</u>	<u>49.1</u>	<u>48.7</u>	<u>32.1</u>
MeWM* [47]	45.2	42.1	43.6	27.9	39.3	<u>48.2</u>	43.3	27.6	41.5	45.3	43.3	27.7
Our Approach	61.3	53.5	57.1	39.9	52.6	48.0	50.2	33.5	56.2	51.8	53.9	36.8

such as Claude-4.5 perform poorly when prompted directly (41.6% F1), our framework successfully leverages the MLLM as part of a simulation-to-decision loop, substantially improving its effective performance. A critical comparison is with MeWM. Direct comparison with the original MeWM [47] was not feasible, as their in-house data is not publicly available and their work focuses on CT data, whereas ours targets multi-sequence MRI. To provide a fair baseline, we re-implemented a MeWM-style diffusion predictor (denoted as MeWM*) using our MRI data. As shown in Tab. 1, our approach (57.1% F1) significantly outperforms this MeWM* baseline (43.6% F1). This result validates our design choice to focus on **time-aware latent dynamics and survival-aware consistency** rather than diffusion-based reconstruction.

Clinical prognostication itself is highly uncertain and challenging, even for experienced practitioners. Prior studies report that neurosurgeons achieve roughly 40% accuracy when predicting outcomes for high-grade glioma patients, while multidisciplinary oncology teams reach approximately 55% accuracy in advanced cancer cohorts [17, 34]. These findings highlight the intrinsic difficulty of survival-oriented treatment planning. CLARITY achieves consistently higher predictive accuracy on the evaluated datasets, suggesting that trajectory-based simulation models can provide valuable decision-support signals for treatment planning.

4.3 Survival Analysis

As shown in Tab. 2, CLARITY significantly outperforms standard survival modeling baselines (DeepSurv [23], RSF [22], SurvTRACE [45]), achieving a state-of-the-art C-index of 0.794.

Beyond quantitative metrics, Fig. 5 demonstrates our model’s superior discriminative capacity in Kaplan-Meier analysis compared to the diffusion-based MeWM baseline. CLARITY yields earlier and

Table 2: C-index performance of various methods.

Method	C-index
DeepSurv	0.664
RSF	0.672
SurvTRACE	0.713
MeWM*	0.701
CLARITY	0.794±0.02




Fig. 5: Kaplan–Meier survival curves predicted by MeWM (left) and our method (right) on MU-Glioma-Post. Our approach produces a much clearer separation across risk strata, reflected by a lower log-rank p -value of **0.0017** and a substantially higher C-index of **0.7943**. Shaded regions denote 95% confidence intervals.

more sustained separation between high- and low-risk strata with minimal curve crossing, resulting in a substantially lower log-rank p -value (0.0017 vs. 0.0763). We attribute this robust stratification to our structured representation space, which avoids diffusion-induced artifacts and preserves structural consistency for reliable forecasting.

4.4 Analysis of Disease Decision Trajectory

Fig. 6 visualizes a predicted multi-stage decision trajectory. Each stage S_i corresponds to an MRI observation and the associated latent state. At every stage, the model performs policy-conditioned state projection, generating multiple candidate branches (dashed lines) corresponding to different treatment actions. The solid lane indicates the selected branch that achieves the lowest predicted risk score, with the corresponding therapy action labeled at each step. Through iterative selection, the model constructs a temporally consistent treatment lane (e.g., RT+TMZ $\rightarrow$ TMZ $\rightarrow$ CCNU $\rightarrow$ Avastin+Brachy) that minimizes the longitudinal risk trajectory.

4.5 Human Evaluation

For therapy recommendation, we use the recorded physician-administered regimen in MU-Glioma-Post as the reference regimen. We report exact-match accuracy against the recorded regimen, while noting that the administered therapy represents an observed clinical decision rather than a unique optimal treatment. To complement this automatic metric, we further conducted a blinded expert assessment on 40 held-out cases. Two medical experts were presented with the patient context and two candidate regimens, corresponding to CLARITY’s recommendation and the physician-administered regimen, with the source labels hidden. The experts selected the preferred regimen or marked the two as clinically equivalent. As shown in Table 3, CLARITY was judged equivalent/non-inferior in 60.0% of cases and preferred over the physician regimen in 35.0% of




Fig. 6: Simulated Multi-Stage Decision Trajectories. Stages S_i denote MRI-anchored follow-ups. Dashed lines show **entropy-regularized exploration** of candidate therapies. The solid path highlights the **optimal sequence** minimizing cumulative risk J , adapting from standard-of-care (*e.g.*, TMZ) to salvage regimens (*e.g.*, CCNU, Avastin) based on state transitions. Boxed numbers denote 30-day risk scores.

Table 3: Blinded expert assessment comparing CLARITY recommendations with recorded physician-administered regimens on 40 held-out MU-Glioma-Post cases.

Outcome	Cases (%)
Equivalent / non-inferior	60.0
CLARITY preferred	35.0
Physician preferred	5.0

cases, while the recorded regimen was preferred in only 5.0% of cases. These results indicate that CLARITY’s recommendations are clinically acceptable in the vast majority of evaluated cases, with 95.0% rated as either equivalent/non-inferior or preferable to the recorded regimen. This suggests that CLARITY does not simply imitate historical physician decisions, but can generate plausible alternative regimens that remain aligned with expert clinical reasoning. By combining guideline-informed candidate generation with survival-conditioned evaluation through the Actor, CLARITY produces patient-specific and outcome-aware therapy recommendations.

4.6 Ablation Study

Impact of CLARITY’s Module. Tab. 4 analyzes the contribution of each component. (1) **Latent vs. Diffusion:** Replacing diffusion synthesis (# 1) with latent dynamics (# 2) yields the largest gain (+8.8% in F1), confirming that compact manifolds capture predictive structure better than stochastic pixel reconstruction. (2) **Clinical Context:** Integrating patient-specific priors (# 3) enhances Recall, indicating that biological signals guide physiologically faithful simulations. (3) **Inverse Evaluation:** The feedback loop (# 4 & 5) drives performance to a peak F1-score of 57.1% (# 5). This validates that iteratively refining therapies via survival feedback is essential for optimized decision-making.

Table 4: Comprehensive Ablation Studies on MU-Glioma-Post. (Top) Component analysis comparing architecture, context, and feedback. (Bottom Left) Impact of iteration K . (Bottom Right) Impact of loss combinations on the C-index.

#	Diff. -based	Latent -based	Context	Feedback Iteration	MU-Glioma-Post (%)		
					Prec.	Rec.	F1
1	✓	✗	✗	✗	47.8	39.8	43.6
2	✗	✓	✗	✗	59.5	46.5	52.4
3	✗	✓	✓	✗	58.5	48.1	52.8
4	✗	✓	✗	✓ (K=3)	59.3	49.6	54.0
5	✗	✓	✓	✓ (K=3)	61.3	53.5	57.1

Iteration Number (K)				Loss Combination	
K	Prec.	Rec.	F1	Loss Components	C-index (%)
1	59.7	48.1	53.1	$\mathcal{L}_{\text{latent}} + \mathcal{L}_{\text{Cox}}$	76.1
2	59.3	48.6	53.4	$\mathcal{L}_{\text{latent}} + \mathcal{L}_{\text{Brier}} + \mathcal{L}_{\text{Cox}}$	76.8
3	61.3	53.5	57.1	$\mathcal{L}_{\text{latent}} + \mathcal{L}_{\text{con}}(\text{Hard}) + \mathcal{L}_{\text{Brier}} + \mathcal{L}_{\text{Cox}}$	77.3
4	57.2	53.6	55.4	$\mathcal{L}_{\text{latent}} + \mathcal{L}_{\text{con}}(\text{Soft}) + \mathcal{L}_{\text{Brier}} + \mathcal{L}_{\text{Cox}}$	79.4

Impact of Iteration Number. Tab. 4 investigates the impact of iteration number (K) on Inverse Survival Evaluation. We observe a steady performance gain as iterations increase, peaking at $K=3$ with a best F1-score of 57.1%. This trajectory validates that iterative feedback enables our method to refine proposed therapies based on accumulated survival signals. However, the slight decline at $K=4$ indicates diminishing returns, suggesting that excessive iterations may introduce noise. Consequently, we adopt $K=3$ as the optimal trade-off.

Impact of Loss Combination. Tab. 4 proves that richer supervision enhances predictive capacity. Adding the Brier score (# 2) improves over the baseline (# 1) by enforcing probability calibration (+0.7%). Incorporating contrastive learning (# 3) further sharpens latent transitions. Crucially, the soft-label variant (# 4) yields the highest C-index (79.4%), confirming that capturing nuanced treatment similarities stabilizes latent dynamics better than hard labels.

Impact of Encoders. As shown in Tab. 5, the choice of visual encoder substantially influences survival prediction performance. Encoders pretrained on natural images (e.g., DINOv2) exhibit limited transferability to longitudinal MRI modeling, yielding the lowest C-index (0.665). In contrast, medical-domain pretraining (MedSigLIP) improves both ranking accuracy and probability calibration, suggesting that domain-aligned representations better capture clinically relevant patterns. Further gains are observed when leveraging MRI-specific encoders (brainIAC [40] and MRI-CORE [14]), which consistently outperform natural and generic medical models. In particular, MRI-CORE achieves the best C-index (0.794) and lowest Brier score (0.162), indicating that modality-aligned pretraining combined with lightweight adaptation (LoRA) provides more discriminative and stable latent representations for downstream survival reasoning.

Comparison between Diffusion and Latent Representation. Tab. 6 highlights the dual advantages of our latent-based dynamics in both computational efficiency and feature fidelity. First, the diffusion-based approach is prohibitively

Table 5: All methods share the same world model and training protocol. We report mean $\pm$ std over 5 random seeds. Higher C-index are better; lower Brier score is better.

Encoder	Pretrain	Frozen	C-index	Brier
DINOv2	Natural (SSL)	Yes	0.665 ± 0.020	0.198 ± 0.008
MedSigLIP	Medical	Yes	0.705 ± 0.016	0.176 ± 0.006
brainIAC	MRI	LoRA	0.786 ± 0.012	0.168 ± 0.004
MRI-CORE	MRI	LoRA	0.794 ± 0.010	0.162 ± 0.003

Table 6: Comparison of efficiency and drift between diffusion and latent methods.

Method	Efficiency		Representation Drift	
	FLOPs (T) $\downarrow$	Time (s) $\downarrow$	Cosine Dist. $\downarrow$	Ratio $\uparrow$
Diffusion-based (1000 steps)	61.3	38.6	0.0174	-
Diffusion-based (500 steps)	39.5	19.7		
Latent-based (Ours)	4.21	0.341	0.0084	1.9$\times$

expensive, requiring up to 61.3 TFLOPs and 38.6 seconds for a *single* simulation. In contrast, our latent-based predictor reduces computation to 4.21 TFLOPs (an approximate 9–15 $\times$ reduction) and achieves sub-second inference (0.341s). This efficiency gain is essential, making our Inverse Survival Evaluation computationally feasible even in a multi-iteration setting (*i.e.*, $K = 3$). Second, we quantify the fidelity loss incurred by the diffusion model’s pixel-space detour (generation followed by re-encoding). Keeping the encoder and survival head identical to ensure a fair comparison, we measure the cosine distance to the ground-truth post-treatment latent. As shown in Tab. 6, diffusion-based reconstruction leads to nearly 2 $\times$ higher representation drift. This indicates that the intermediate pixel-generation step introduces substantial noise, whereas our direct latent prediction maintains the stability required for downstream survival evaluation.

5 Conclusion

We introduced CLARITY, a medical world model addressing critical limitations in prior works. By forecasting disease evolution in a **latent space**, we avoid the stochasticity and high computational cost of diffusion-based models. CLARITY is the first to explicitly integrate **temporal and clinical contexts** with an **Inverse Survival Evaluation**. Experiment proves that our feedback loop is critical for optimized decisions, with results outperforming all baselines. Therefore, CLARITY represents a significant step towards computationally feasible and personalized treatment planning in oncology.

References

1. Al Radi, A.M., Cao, X., Yu, F., Liu, Y., Liu, F., Wang, C., Chen, Y., Chen, J., Wang, H., Meng, Y., et al.: Agentic large-language-model systems in medicine: A systematic review and taxonomy. Authorea Preprints (2025)
2. Anthropic, A.: Claude 3.5 sonnet model card addendum. Claude-3.5 Model Card 3(6) (2024)

3. Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., Ballas, N.: Self-supervised learning from images with a joint-embedding predictive architecture. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15619–15629 (2023)
4. Assran, M., Bardes, A., Fan, D., Garrido, Q., Howes, R., Muckley, M., Rizvi, A., Roberts, C., Sinha, K., Zholus, A., et al.: V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985* (2025)
5. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923* (2025)
6. Bardes, A., Garrido, Q., Ponce, J., Chen, X., Rabbat, M., LeCun, Y., Assran, M., Ballas, N.: Revisiting feature prediction for learning visual representations from video. *arXiv preprint arXiv:2404.08471* (2024)
7. Bruce, J., Dennis, M.D., Edwards, A., Parker-Holder, J., Shi, Y., Hughes, E., Lai, M., Mavalankar, A., Steigerwald, R., Apps, C., et al.: Genie: Generative interactive environments. In: *Forty-first International Conference on Machine Learning* (2024)
8. Burchi, M., Timofte, R.: Mudreamer: Learning predictive world models without reconstruction. *arXiv preprint arXiv:2405.15083* (2024)
9. Chen, J., Gui, C., Ouyang, R., Gao, A., Chen, S., Chen, G.H., Wang, X., Zhang, R., Cai, Z., Ji, K., et al.: Huatuogpt-vision, towards injecting medical visual knowledge into multimodal llms at scale. *arXiv preprint arXiv:2406.19280* (2024)
10. Chen, Y., Liu, Y., Wang, C., Elliott, M., Kwok, C.F., Peña-Solorzano, C., Tian, Y., Liu, F., Frazer, H., McCarthy, D.J., et al.: Braixdet: Learning to detect malignant breast lesion with incomplete annotations. *Medical image analysis* **96**, 103192 (2024)
11. Cheng, P., Lin, L., Lyu, J., Huang, Y., Luo, W., Tang, X.: Prior: Prototype representation joint learning from medical images and reports. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 21361–21371 (October 2023)
12. Cox, D.R.: Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* **34**(2), 187–202 (1972)
13. Doerr, A., Daniel, C., Schiegg, M., Duy, N.T., Schaal, S., Toussaint, M., Sebastian, T.: Probabilistic recurrent state-space models. In: *International conference on machine learning*. pp. 1280–1289. PMLR (2018)
14. Dong, H., Chen, Y., Gu, H., Konz, N., Chen, Y., Li, Q., Mazurowski, M.A.: Mri-core: a foundation model for magnetic resonance imaging. *arXiv preprint arXiv:2506.12186* (2025)
15. Fields, B.K., Calabrese, E., Mongan, J., Cha, S., Hess, C.P., Sugrue, L.P., Chang, S.M., Luks, T.L., Villanueva-Meyer, J.E., Rauschecker, A.M., et al.: The university of california san francisco adult longitudinal post-treatment diffuse glioma mri dataset. *Radiology: Artificial Intelligence* **6**(4), e230182 (2024)
16. Goma, A., Huang, Y., Hagag, A., Schmitter, C., Höfler, D., Weissmann, T., Breininger, K., Schmidt, M., Stritzelberger, J., Delev, D., et al.: Comprehensive multimodal deep learning survival prediction enabled by a transformer architecture: A multicenter study in glioblastoma. *Neuro-Oncology Advances* **6**(1), vdae122 (2024)
17. Gwilliam, B., Keeley, V., Todd, C., Roberts, C., Gittins, M., Kelly, L., Barclay, S., Stone, P.: Prognosticating in patients with advanced cancer—observational study comparing the accuracy of clinicians’ and patients’ estimates of survival. *Annals of oncology* **24**(2), 482–488 (2013)

18. Hafner, D., Pasukonis, J., Ba, J., Lillicrap, T.: Mastering diverse control tasks through world models. *Nature* pp. 1–7 (2025)
19. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: low-rank adaptation of large language models. *arxiv preprint arXiv:2106.09685* (2021)
20. Huang, S.C., Shen, L., Lungren, M.P., Yeung, S.: Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3942–3951 (2021)
21. Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. *arXiv preprint arXiv:2410.21276* (2024)
22. Ishwaran, H., Kogalur, U.B., Blackstone, E.H., Lauer, M.S.: *Random survival forests* (2008)
23. Katzman, J.L., Shaham, U., Cloninger, A., Bates, J., Jiang, T., Kluger, Y.: Deep-surv: personalized treatment recommender system using a cox proportional hazards deep neural network. *BMC medical research methodology* **18**(1), 24 (2018)
24. Koju, S., Bastola, S., Shrestha, P., Amgain, S., Shrestha, Y.R., Poudel, R.P., Bhattarai, B.: Surgical vision world model. In: *MICCAI Workshop on Data Engineering in Medical Imaging*. pp. 1–10. Springer (2025)
25. Kraljevic, Z., Bean, D., Shek, A., Bendayan, R., Hemingway, H., Yeung, J.A., Deng, A., Balston, A., Ross, J., Idowu, E., et al.: Foresight—a generative pretrained transformer for modelling of patient timelines using electronic health records: a retrospective modelling study. *The Lancet Digital Health* **6**(4), e281–e290 (2024)
26. Lee, C., Yoon, J., Van Der Schaar, M.: Dynamic-deephit: A deep learning approach for dynamic survival analysis with competing risks based on longitudinal data. *IEEE Transactions on Biomedical Engineering* **67**(1), 122–133 (2019)
27. Lee, C., Zame, W., Yoon, J., Van Der Schaar, M.: Deephit: A deep learning approach to survival analysis with competing risks. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 32 (2018)
28. Li, M., Wen, C., Tian, Y., Shi, M., Luo, Y., Huang, H., Fang, Y., Wang, M.: Fairfedmed: Benchmarking group fairness in federated medical imaging with fair-lora. *IEEE Transactions on Medical Imaging* (2025)
29. Liu, Y., Tian, Y., Wang, C., Chen, Y., Liu, F., Belagiannis, V., Carneiro, G.: Translation consistent semi-supervised segmentation for 3d medical images. *IEEE Transactions on Medical Imaging* (2024)
30. Luo, Y., Shi, M., Tian, Y., Elze, T., Wang, M.: Harvard glaucoma detection and progression: A multimodal multitask dataset and generalization-reinforced semi-supervised learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 20471–20482 (2023)
31. Mahmoudi, K., Kim, D.H., Tavakkol, E., Kihira, S., Bauer, A., Tsankova, N., Khan, F., Hormigo, A., Yedavalli, V., Nael, K.: Multiparametric radiogenomic model to predict survival in patients with glioblastoma. *Cancers* **16**(3), 589 (2024)
32. Moon, I., Groha, S., Gusev, A.: Survlatent ode: A neural ode based time-to-event model with competing risks for longitudinal data improves cancer-associated venous thromboembolism (vte) prediction. In: *Machine Learning for Healthcare Conference*. pp. 800–827. PMLR (2022)
33. Nagpal, C., Li, X., Dubrawski, A.: Deep survival machines: Fully parametric survival regression and representation learning for censored data with competing risks. *IEEE Journal of Biomedical and Health Informatics* **25**(8), 3163–3175 (2021)

34. Sagberg, L.M., Jakola, A.S., Reinertsen, I., Solheim, O.: How well do neurosurgeons predict survival in patients with high-grade glioma? *Neurosurgical Review* **45**(1), 865–872 (2022)
35. Sellergren, A., Kazemzadeh, S., Jaroensri, T., Kiraly, A., Traverse, M., Kohlberger, T., Xu, S., Jamil, F., Hughes, C., Lau, C., et al.: Medgemma technical report. arXiv preprint arXiv:2507.05201 (2025)
36. Shi, M., Luo, Y., Tian, Y., Shen, L.Q., Zebardast, N., Eslami, M., Kazeminasab, S., Boland, M.V., Friedman, D.S., Pasquale, L.R., et al.: Equitable artificial intelligence for glaucoma screening with fair identity normalization. *NPJ Digital Medicine* **8**(1), 46 (2025)
37. Singh, A., Fry, A., Perelman, A., Tart, A., Ganesh, A., El-Kishky, A., McLaughlin, A., Low, A., Ostrow, A., Ananthram, A., et al.: Openai gpt-5 system card. arXiv preprint arXiv:2601.03267 (2025)
38. Singhal, K., Azizi, S., Tu, T., Mahdavi, S.S., Wei, J., Chung, H.W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., et al.: Large language models encode clinical knowledge. *Nature* **620**(7972), 172–180 (2023)
39. Singhal, K., Tu, T., Gottweis, J., Sayres, R., Wulczyn, E., Amin, M., Hou, L., Clark, K., Pfohl, S.R., Cole-Lewis, H., et al.: Toward expert-level medical question answering with large language models. *Nature Medicine* **31**(3), 943–950 (2025)
40. Tak, D., Gormosa, B., Zapaishchykova, A., et al.: A generalizable foundation model for analysis of human brain mri. *Nature Neuroscience* (2026). <https://doi.org/10.1038/s41593-026-02202-6>, <https://doi.org/10.1038/s41593-026-02202-6>
41. Team, G., Anil, R., Borgeaud, S., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., Millican, K., et al.: Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805 (2023)
42. Team, Q.: Qwen3 technical report (2025), <https://arxiv.org/abs/2505.09388>
43. Wang, F., Zhou, Y., Wang, S., Vardhanabhuti, V., Yu, L.: Multi-granularity cross-modal alignment for generalized medical visual representation learning. *Advances in Neural Information Processing Systems* **35**, 33536–33549 (2022)
44. Wang, H., Yee, D.: I-spy 2: a neoadjuvant adaptive clinical trial designed to improve outcomes in high-risk breast cancer. *Current breast cancer reports* **11**(4), 303–310 (2019)
45. Wang, Z., Sun, J.: Survtrace: Transformers for survival analysis with competing events. In: *Proceedings of the 13th ACM international conference on bioinformatics, computational biology and health informatics*. pp. 1–9 (2022)
46. Xu, M.: Medicalgpt: Training medical gpt model. <https://github.com/shibing624/MedicalGPT> (2023)
47. Yang, Y., Wang, Z.Y., Liu, Q., Sun, S., Wang, K., Chellappa, R., Zhou, Z., Yuille, A., Zhu, L., Zhang, Y.D., et al.: Medical world model: Generative simulation of tumor evolution for treatment planning. arXiv preprint arXiv:2506.02327 (2025)
48. Yaseen, D., Garrett, F., Gass, J., Greaser, J., Isufi, E., Layfield, L., Nada, A., Porgorzelski, K., Sinclair, J., Tahon, N., et al.: University of missouri post-operative glioma dataset (mu-glioma-post)(version 1)[data set]. *The Cancer Imaging Archive* (2025)
49. Yue, Y., Wang, Y., Jiang, H., Liu, P., Song, S., Huang, G.: Echoworld: Learning motion-aware world models for echocardiography probe guidance. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 25993–26003 (2025)
50. Yue, Y., Wang, Y., Tao, C., Liu, P., Song, S., Huang, G.: Chexworld: Exploring image world modeling for radiograph representation learning. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 20778–20788 (2025)

51. Zou, Y., Yin, Z.: Alignment, mining and fusion: Representation alignment with hard negative mining and selective knowledge fusion for medical visual question answering. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 29623–29633 (2025)
52. Zou, Y., Yin, Z.: Mvcm: Enhancing multi-view and cross-modality alignment for medical visual question answering and medical image-text retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 180–190 (June 2025)