NEUROBELIEF: OBSERVER-CENTRIC WORLD MODELING FOR INFORMATION ACQUISITION IN MULTISEQUENCE MRI

Anonymous authors
Paper under double-blind review

ABSTRACT

World models typically predict how a system changes after an action. In information acquisition, however, the underlying system may remain unchanged while the observer's task state changes as new evidence becomes available. This raises a different prediction problem: before acquiring a candidate observation, can we anticipate how it would update the current task state? We study this problem in multisequence brain MRI and introduce NeuroBelief. The currently observed sequences are encoded into a spatial Pathology Belief. Given this belief and the identity of a candidate sequence, NeuroBelief predicts the Pathology Belief that would be formed after acquisition, without access to the candidate image. Once the sequence is acquired, the same encoder constructs the corresponding target state, allowing the predicted transition to be evaluated directly. On BraTS 2023 and BraTS Africa, the predicted transitions depend on both the current patient state and the candidate sequence, and generalize to observation and acquisition combinations excluded from transition training. Using these predicted states for sequential acquisition, NeuroBelief reduces the number of acquired sequences by 48.4% internally and 44.7% externally relative to acquiring all four sequences, while mean Dice remains 0.68 and 1.41 percentage points below the corresponding four sequence references. These results show that world modeling can also be used to predict how new observations change an observer's task state, providing a directly testable basis for adaptive information acquisition.

1 INTRODUCTION

World models support reasoning, planning, and control by predicting how a system changes after an action (Hafner et al., 2019; 2020; Hansen et al., 2024). In medical examination, active perception, and scientific measurement, many actions instead acquire new evidence. This new evidence updates the task state maintained by the observer. This leads to a central question for information acquisition: *before new evidence becomes available, can a model predict the task state that each candidate observation would produce?*

We formulate this setting as *world modeling from the observer's perspective*. Given the current evidence and a candidate observation, the model predicts the task state that would result from acquiring that observation. The prediction is made before the candidate observation itself becomes available. Once the observation is acquired, the enlarged evidence set is encoded into the corresponding target state. The predicted and observed states can therefore be compared directly (Fig. 1). This formulation separates two capabilities. Action selection determines which candidate observation to acquire, whereas state prediction describes how each candidate observation is expected to update the current task representation.

Multisequence MRI provides a structured and practically relevant setting for studying this problem. Glioma assessment commonly combines native T1 weighted, contrast enhanced T1 weighted, T2 weighted, and T2 FLAIR sequences because they provide complementary information about tumor extent and tissue characteristics (Weller et al., 2021; Wen et al., 2023). Clinical MRI for brain tumors commonly follows a predefined protocol with a fixed set of sequences. A basic examination typically requires approximately 25–30 minutes (Thust et al., 2018). However, the additional task

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

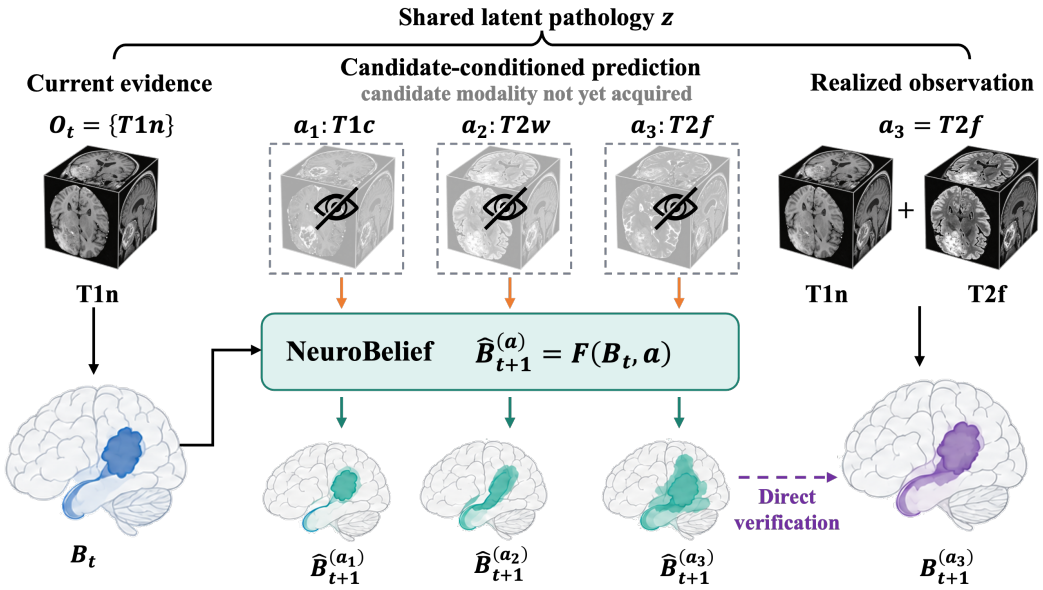

Figure 1: Predictive state modeling from the observer's perspective. Candidate state changes are predicted before acquisition and verified afterward.

information provided by a new sequence may differ across patients and across observation states. This motivates acquisition that adapts to the available evidence. After each acquired sequence, the model updates the current patient state and predicts the Pathology Belief that would result from acquiring each remaining candidate sequence. These predicted consequences are then used to select the next sequence or stop acquisition while balancing task performance against the number of acquired sequences.

Existing studies of incomplete MRI observations primarily investigate how downstream tasks such as segmentation can be performed when some sequences are unavailable (Konwer et al., 2023; Wang et al., 2023; Liu et al., 2023; Yang et al., 2023; Qiu et al., 2025; Li et al., 2025; Zhu et al., 2026). Methods for active information acquisition instead focus on estimating the utility of candidate observations or learning which observation to acquire next (Ma et al., 2019; Valancius et al., 2024; Norcliffe et al., 2025; Shim et al., 2018; Guney et al., 2025; Li & Oliva, 2021). These two lines of work address different questions: what can be inferred from the evidence already available, and what should be observed next? We study the prediction problem between them. Before acquisition, we predict the task state that a candidate observation would produce. After the observation becomes available, that prediction can be tested directly.

We introduce NeuroBelief to model how candidate observations change the task state. The model first encodes the available MRI sequences into a spatial *Pathology Belief*. It then predicts how this state would change for each candidate sequence using only the current belief and candidate identity. The candidate image itself is not available to the transition model. Once a candidate sequence is acquired, the same encoder constructs the corresponding target state from the enlarged observation set. Predicted and target states lie in the same representation space and are grounded in the segmentation task, so the predicted update can be evaluated directly. The predicted states can also support subsequent acquisition decisions.

Our main contributions are:

- We formulate *observer-centric world modeling* for information acquisition, where an acquisition action is represented by its effect on the observer's task state rather than the underlying physical system. This makes the consequence of a candidate observation predictable before acquisition and directly verifiable afterward.
- We introduce NeuroBelief, which predicts the spatial Pathology Belief associated with each candidate sequence from the current patient state and candidate identity, without access to

the candidate image. The predicted transitions capture both patient and candidate dependence and generalize to observation and acquisition combinations excluded from transition training.

- We demonstrate that the predicted states can support sequential MRI acquisition, enabling adaptive acquisition for individual patients while reducing the number of acquired sequences and maintaining segmentation performance close to the four sequence reference.

2 Related Work

World models. World models predict how a system changes after an action and use these predictions for planning and control (Hafner et al., 2019; 2020; Hansen et al., 2024). Medical world models have applied this idea to disease evolution and imaging state changes, including tumor progression under treatment and spatial changes caused by probe movement in ultrasound (Yang et al., 2025; Jiang et al., 2024; Kong et al., 2026; Yang et al., 2026). NeuroBelief considers information acquisition instead. Given the observed MRI sequences and a candidate sequence identity, it predicts the spatial Pathology Belief that would be formed after acquisition without accessing the candidate image.

Belief state and predictive state models. Belief state planning in partially observable Markov decision processes represents incomplete information as a posterior distribution over latent states (Kaelbling et al., 1998). Predictive state representations describe a system through predictions of future observations under possible actions (Littman et al., 2001). NeuroBelief instead uses a spatial task state learned with direct supervision. For each candidate observation, it predicts the state after acquisition and compares it with the state encoded once the observation becomes available.

Learning from incomplete observations and active information acquisition. Methods for incomplete MRI observations aim to preserve downstream performance when some sequences are unavailable (Konwer et al., 2023; Wang et al., 2023; Liu et al., 2023; Yang et al., 2023; Qiu et al., 2025; Li et al., 2025; Zhu et al., 2026). Active information acquisition instead focuses on deciding which additional observation should be obtained next (Ma et al., 2019; Valancius et al., 2024; Norcliffe et al., 2025; Shim et al., 2018; Guney et al., 2025; Li & Oliva, 2021). Adaptive breast MRI has also studied whether additional sequences are needed after an initial set of images (Eskreis-Winkler et al., 2025).

NeuroBelief addresses the prediction problem between these settings. Rather than only inferring from current observations or directly selecting the next acquisition, it predicts the spatial Pathology Belief associated with each candidate sequence before acquisition. These predicted states then guide subsequent MRI acquisition.

3 NeuroBelief

NeuroBelief models how the task state maintained by an observer changes as additional MRI evidence becomes available. As shown in Fig. 2, NeuroBelief first constructs a spatial task state from the currently observed sequences. It then predicts the state associated with each candidate acquisition without access to the candidate image and evaluates the predicted states for subsequent sequence selection.

3.1 Predictive States from the Observer Perspective

A conventional world model predicts how a system state changes after an action:

$$s_{t+1} = f(s_t, a_t). \tag{1}$$

In NeuroBelief, an action acquires information rather than changing the physical system. We treat the underlying pathology as unchanged within one acquisition step. The action instead changes the evidence available to the observer and therefore may change the task state supported by that evidence.

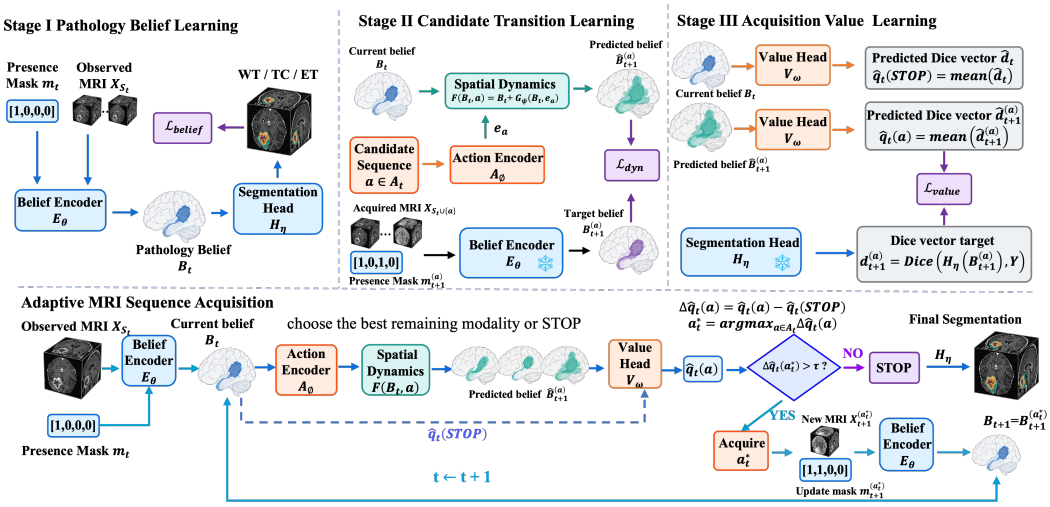

Figure 2: Overview of NeuroBelief. Predicted candidate states are evaluated to select the next sequence or stop acquisition.

Let $\mathcal{M}$ denote the set of available MRI sequence types. At step t, $S_t \subseteq \mathcal{M}$ contains the observed sequences, and X_{S_t} denotes their images. The remaining candidate set is

$$\mathcal{A}_t = \mathcal{M} \setminus S_t. \tag{2}$$

The current observer state is

$$B_t = E_\theta(X_{S_t}, m_t), \tag{3}$$

where m_t indicates which sequences are present.

For a candidate $a \in \mathcal{A}_t$, we predict the state associated with acquiring that sequence:

$$\widehat{B}^{(a)}_{t+1} = F(B_t, a). \tag{4}$$

This prediction uses the current observer state and the identity of the candidate sequence. The candidate image itself is not available to the transition model.

After acquisition, the candidate image is added to the observed set, and the same encoder constructs the corresponding target state:

$$B^{(a)}_{t+1} = E_\theta\left(X_{S_t \cup \{a\}}, m^{(a)}_{t+1}\right). \tag{5}$$

The predicted state $\widehat{B}^{(a)}_{t+1}$ and the target state $B^{(a)}_{t+1}$ lie in the same representation space. The consequence of a candidate acquisition can therefore be predicted before the image is observed and evaluated after it becomes available. NeuroBelief predicts the resulting task state rather than reconstructing the unseen MRI image.

3.2 MRI Observer State and Candidate Transition

At each acquisition step, the model receives four MRI channels together with a presence mask. Missing sequences are filled with zeros. The Belief Encoder constructs

$$B_t = E_\theta(X_{S_t}, m_t), \qquad B_t \in \mathbb{R}^{C_B \times D_B \times H_B \times W_B}. \tag{6}$$

We call B_t the *Pathology Belief*. It is a deterministic spatial representation of the task state supported by the currently available MRI evidence. The Belief Encoder is trained with WT, TC, and ET segmentation supervision so that this representation remains grounded in tumor segmentation. Here, "belief" refers to the task state inferred from the available observations rather than a probabilistic posterior over pathology.

For each candidate $a \in \mathcal{A}_t$, the Action Encoder maps its sequence identity to

$$e_a = A_\phi(a). \tag{7}$$

Spatial Dynamics then predicts the Pathology Belief associated with acquiring that candidate through a residual update:

$$\widehat{B}_{t+1}^{(a)} = F(B_t, a) = B_t + G_\psi(B_t, e_a). \tag{8}$$

During Dynamics training, the Belief Encoder is frozen. The candidate image is used only to construct the target state:

$$B_{t+1}^{(a)} = E_\theta\left(X_{S_t \cup \{a\}}, m_{t+1}^{(a)}\right). \tag{9}$$

Spatial Dynamics is therefore trained to predict the Pathology Belief after acquisition from B_t and a alone. The candidate image determines the supervision target but is never provided as input to the transition model.

3.3 Adaptive Acquisition from Predicted States

NeuroBelief estimates the segmentation utility associated with the current state and each predicted future state. The Value Head predicts Dice scores for WT, TC, and ET:

$$\hat{\mathbf{d}}_t = V_\omega(B_t), \qquad \hat{\mathbf{d}}_{t+1}^{(a)} = V_\omega\left(\widehat{B}_{t+1}^{(a)}\right). \tag{10}$$

The corresponding training targets are computed using the frozen Segmentation Head and the segmentation annotation:

$$\mathbf{d}_t = \mathrm{Dice}\left(H_\eta(B_t), Y\right), \qquad \mathbf{d}_{t+1}^{(a)} = \mathrm{Dice}\left(H_\eta(B_{t+1}^{(a)}), Y\right), \tag{11}$$

where Y is the segmentation annotation.

The mean predicted Dice is used as the task value of the current state and each candidate acquisition:

$$\hat{q}_t(\mathrm{STOP}) = \mathrm{Mean}\left(\hat{\mathbf{d}}_t\right), \qquad \hat{q}_t(a) = \mathrm{Mean}\left(\hat{\mathbf{d}}_{t+1}^{(a)}\right). \tag{12}$$

Here, $\hat{q}_t(\mathrm{STOP})$ is the value of retaining the current observation set, while $\hat{q}_t(a)$ is the predicted value after acquiring candidate a. Throughout this work, acquisition value refers to expected segmentation performance under the frozen task model rather than clinical utility or information gain.

The predicted gain of candidate a is

$$\Delta\hat{q}_t(a) = \hat{q}_t(a) - \hat{q}_t(\mathrm{STOP}). \tag{13}$$

The next sequence is selected by

$$a_t^* = \arg\max_{a \in \mathcal{A}_t} \Delta\hat{q}_t(a). \tag{14}$$

Acquisition stops when

$$\Delta\hat{q}_t(a_t^*) \le \tau. \tag{15}$$

Otherwise, sequence a_t^* is added to the observed set. The Pathology Belief is recomputed from the enlarged set of acquired sequences, and the remaining candidates are evaluated from the updated state. The final segmentation uses only the sequences acquired before stopping.

3.4 Learning Objectives

NeuroBelief is trained in three stages. Stage I learns the Pathology Belief from incomplete observations. Stage II freezes the Belief Encoder and learns candidate transitions. Stage III freezes the preceding components and trains the Value Head for task value and candidate ranking. Complete objectives and loss definitions are provided in Appendix A.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

We use BraTS 2023 for model development and internal evaluation, with 1,001 training, 125 development, and 125 test cases with no patient overlap across the three sets (BraTS 2023 Challenge, 2023). External evaluation is performed on 95 BraTS-Africa cases without retraining or adaptation (Adewole et al., 2025). Model selection, hyperparameter tuning, and policy development use only the training and development sets. Detailed preprocessing, optimization settings, and training configurations are provided in Appendix B.

4.2 EVALUATION PROTOCOL

We evaluate NeuroBelief at three levels. First, *consequence prediction* measures whether the state predicted before acquisition matches the state obtained after the candidate sequence becomes available. We further test whether these predictions use information from both the current patient state and the candidate identity. Second, *transition generalization* evaluates observation and acquisition combinations that are excluded from Dynamics training, testing whether the learned transition structure transfers beyond the combinations seen during optimization. Third, *sequential acquisition* evaluates the predicted states in closed loop, where the model repeatedly selects a remaining sequence or stops based on the updated patient state. Performance is assessed jointly by the final segmentation accuracy and the number of acquired sequences.

5 RESULTS

5.1 CANDIDATE MODALITY EFFECTS DEPEND ON THE PATIENT STATE

We first test whether NeuroBelief can predict the state change caused by an unobserved candidate modality. Identity assumes no change, Spatial Population uses the mean training update for the same observation subset and candidate modality, and Direct Logit Delta predicts the segmentation output change directly without constructing a future Pathology Belief. Detailed definitions and evaluation procedures are provided in Appendix C.

As shown in Table 1, NeuroBelief consistently improves over Identity and Spatial Population on both cohorts. Direct Logit Delta achieves similar spatial prediction accuracy, but NeuroBelief better preserves the relative task consequences of the available candidates, with lower Gain MAE and regret and higher Top 1 accuracy. Replacing the current belief with another patient's state increases Change MAE from 4.718 to 5.799 internally and from 11.045 to 12.314 externally. Permuting candidate identity reduces centered candidate identification from 0.817 to 0.119 and from 0.768 to 0.166, respectively. These results show that the predicted transition depends on both the current patient state and the candidate modality. Fig. 3 shows representative spatial changes predicted before candidate acquisition.

5.2 NEUROBELIEF GENERALIZES TO WITHHELD TRANSITION COMBINATIONS

We next test whether Spatial Dynamics can generalize to transition combinations that are absent from training. The 24 valid transitions with one or two observed modalities are divided into three folds. For each fold, eight transitions are removed from Dynamics training and used only for evaluation. Seen Spatial Population estimates the mean update from available training transitions with the same observation size and candidate modality. The complete protocol is provided in Appendix D.

NeuroBelief consistently outperforms Seen Spatial Population on the withheld transitions (Table 2). Latent L1 decreases from 0.159 ± 0.003 to 0.111 ± 0.004 on BraTS 2023 and from 0.164 ± 0.006 to 0.118 ± 0.007 on BraTS Africa. The decoded change metrics and candidate identification show the same pattern. These results indicate that Spatial Dynamics can predict state changes for observation and acquisition combinations that were not included in transition training.

Table 1: Prediction of spatial consequences (A) and task consequences (B). Values are patient means ± sample standard deviations after averaging matched transitions within each patient.

A. Spatial consequence				
Cohort	**Predictor**	**Change MAE** ($\times 0.001$) ↓	**Pearson** ↑	**Top 10% Dice** ↑
BraTS 2023	Identity	5.330 ± 3.176	–	–
	Spatial Population	5.195 ± 2.946	0.556 ± 0.038	0.871 ± 0.029
	Direct Logit Delta	4.724 ± 2.874	0.580 ± 0.043	0.876 ± 0.030
	NeuroBelief	$\mathbf{4.718 \pm 2.853}$	$\mathbf{0.581 \pm 0.033}$	$\mathbf{0.885 \pm 0.030}$
BraTS Africa	Identity	11.864 ± 7.398	–	–
	Spatial Population	11.633 ± 7.264	0.545 ± 0.036	0.847 ± 0.058
	Direct Logit Delta	11.206 ± 7.150	0.551 ± 0.040	0.850 ± 0.059
	NeuroBelief	$\mathbf{11.045 \pm 7.055}$	$\mathbf{0.560 \pm 0.036}$	$\mathbf{0.854 \pm 0.060}$
B. Task consequence				
Cohort	**Predictor**	**Gain MAE** ↓	**Top 1** ↑	**Regret** ↓
BraTS 2023	Direct Logit Delta	0.064 ± 0.030	0.488 ± 0.252	0.038 ± 0.035
	NeuroBelief	$\mathbf{0.054 \pm 0.029}$	$\mathbf{0.602 \pm 0.179}$	$\mathbf{0.024 \pm 0.028}$
BraTS Africa	Direct Logit Delta	0.101 ± 0.040	0.336 ± 0.209	0.069 ± 0.057
	NeuroBelief	$\mathbf{0.083 \pm 0.034}$	$\mathbf{0.477 \pm 0.212}$	$\mathbf{0.050 \pm 0.050}$

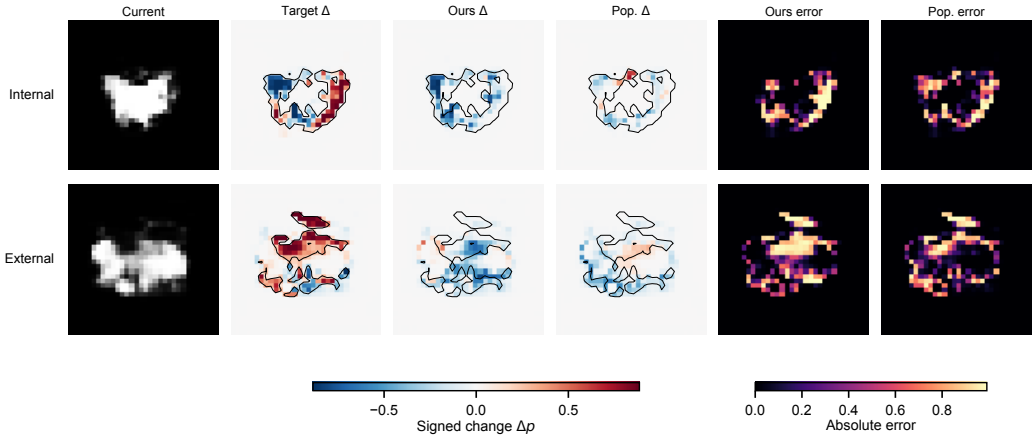

Figure 3: Representative predicted Pathology Belief changes. Contours mark the largest 10% of target change magnitude.

5.3 Task Grounding and Closed Loop Acquisition

We first evaluate whether the Pathology Belief retains task information when only part of the MRI is observed. The frozen NeuroBelief is compared with RFNet (Ding et al., 2021), mmFormer (Zhang et al., 2022), M3AE (Liu et al., 2023), and UniME (Song et al., 2026) using the same patients and observation subsets. Dynamics and the acquisition policy are not used in this comparison. NeuroBelief achieves the highest mean Dice at every missingness level in both cohorts (Table 3, Panel A), with a larger gap as fewer sequences are available.

We then evaluate sequential acquisition from T1c. After each acquisition, the real image is added to the observed set and the Pathology Belief is recomputed before the next decision. At the stopping threshold of 0.03 selected on the development cohort, NeuroBelief reduces sequence acquisition by $48.40 \pm 6.92\%$ on BraTS 2023 and $44.74 \pm 10.25\%$ on BraTS Africa. The corresponding final Dice values are 0.8473 ± 0.1294 and 0.8192 ± 0.1276 (Table 3, Panel B). NeuroBelief is comparable to the population policies internally and achieves the highest mean Dice among the evaluated policies on BraTS Africa. Fig. 5 summarizes the corresponding stopping outcomes. Additional analyses of adaptive opportunity, matched acquisition budgets, and acquisition paths are provided in Appendix E.

Table 2: Values are patient means ± sample standard deviations after averaging evaluated transitions and folds within each patient.

Cohort	Method	Latent L1 ↓	Change MAE ($\times 10^{-3}$) ↓	Direction ↑	Top 10% Dice ↑	Candidate ID ↑
BraTS 2023	Seen Spatial Population	0.159 ± 0.003	5.228 ± 2.948	0.142 ± 0.069	0.855 ± 0.048	0.520 ± 0.095
	NeuroBelief	$\mathbf{0.111 \pm 0.004}$	$\mathbf{4.932 \pm 2.995}$	$\mathbf{0.233 \pm 0.063}$	$\mathbf{0.874 \pm 0.041}$	$\mathbf{0.576 \pm 0.085}$
BraTS Africa	Seen Spatial Population	0.164 ± 0.006	11.673 ± 7.251	0.110 ± 0.068	0.838 ± 0.063	0.498 ± 0.089
	NeuroBelief	$\mathbf{0.118 \pm 0.007}$	$\mathbf{11.477 \pm 7.271}$	$\mathbf{0.177 \pm 0.083}$	$\mathbf{0.847 \pm 0.063}$	$\mathbf{0.541 \pm 0.092}$

Table 3: Task grounding (A) and closed loop acquisition (B). Missingness of 0%, 25%, 50%, and 75% corresponds to 4, 3, 2, and 1 observed sequences. Values are patient means ± sample standard deviation; reduction is relative to four sequences.

A. Fixed observation subsets

BraTS 2023 internal

Method	0%	25%	50%	75%
RFNet	0.693 ± 0.245	0.637 ± 0.227	0.556 ± 0.204	0.400 ± 0.154
mmFormer	0.530 ± 0.268	0.488 ± 0.250	0.422 ± 0.226	0.302 ± 0.177
M3AE	0.706 ± 0.253	0.671 ± 0.242	0.613 ± 0.232	0.534 ± 0.225
UniME	0.827 ± 0.167	0.788 ± 0.172	0.739 ± 0.175	0.646 ± 0.188
NeuroBelief	$\mathbf{0.854 \pm 0.117}$	$\mathbf{0.823 \pm 0.120}$	$\mathbf{0.781 \pm 0.129}$	$\mathbf{0.712 \pm 0.140}$

BraTS Africa external

Method	0%	25%	50%	75%
RFNet	0.714 ± 0.224	0.666 ± 0.209	0.597 ± 0.194	0.466 ± 0.173
mmFormer	0.612 ± 0.263	0.570 ± 0.256	0.504 ± 0.243	0.378 ± 0.204
M3AE	0.683 ± 0.241	0.661 ± 0.232	0.622 ± 0.225	0.564 ± 0.207
UniME	0.814 ± 0.198	0.772 ± 0.192	0.713 ± 0.187	0.604 ± 0.186
NeuroBelief	$\mathbf{0.834 \pm 0.123}$	$\mathbf{0.794 \pm 0.124}$	$\mathbf{0.741 \pm 0.128}$	$\mathbf{0.642 \pm 0.128}$

B. Patient level closed loop acquisition

	BraTS 2023 internal		BraTS Africa external	
Policy	Final Dice ↑	Reduction ↑	Final Dice ↑	Reduction ↑
Population Fixed	$\mathbf{0.8478 \pm 0.1203}$	$50.00 \pm 0.00\%$	0.8085 ± 0.1351	$50.00 \pm 0.00\%$
Spatial Population	0.8458 ± 0.1295	$\mathbf{50.60 \pm 3.84\%}$	0.8047 ± 0.1394	$\mathbf{50.79 \pm 4.40\%}$
NeuroBelief	0.8473 ± 0.1294	$48.40 \pm 6.92\%$	$\mathbf{0.8192 \pm 0.1276}$	$44.74 \pm 10.25\%$

In the representative cases shown in Fig. 4, each newly acquired sequence adds task information to the Pathology Belief about WT, TC, and ET. Some cases approach the four sequence result after one additional sequence and stop, whereas others continue acquisition when further information remains useful. Across the full cohorts, the stopping outcomes in Fig. 5 show that some patients terminate acquisition with fewer than four sequences while remaining within the predefined Dice tolerance of the complete sequence reference. The cumulative coverage indicates how this proportion changes as additional modalities are acquired.

6 Discussion

The central result is that changes in the task state associated with each candidate can be predicted before acquisition and verified afterward. Patient and candidate permutations show that these updates use both conditioning variables, while the transition holdout experiment shows transfer to observation and acquisition combinations withheld from Dynamics training. Direct Logit Delta nearly matched spatial fidelity for individual transitions, whereas NeuroBelief achieved lower gain MAE, higher Top 1 identification accuracy, and lower regret. The main empirical separation therefore occurs in preserving differences in task consequence across candidates.

The relationship between segmentation accuracy and acquisition budget varied across cohorts. NeuroBelief performed comparably to the population policies on the internal cohort and outperformed Population Fixed at several exploratory operating points with matched acquisition budgets on the external cohort. At the primary operating point, however, its Dice remained below the reference obtained with all four sequences in both cohorts. Pathology Belief is a deterministic task representation produced by the encoder and does not quantify uncertainty in candidate outcomes. Future

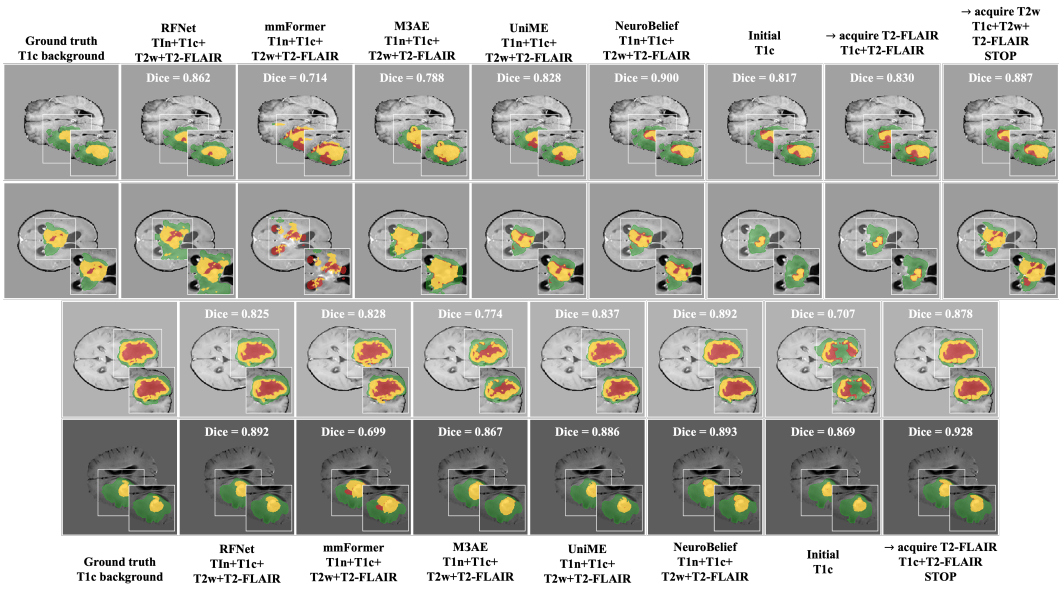

Figure 4: Representative segmentation results and adaptive acquisition trajectories. Rows denote patients; columns show the reference, four sequence predictions, and NeuroBelief acquisition steps. Modality sets and Dice scores are shown above and below each prediction.

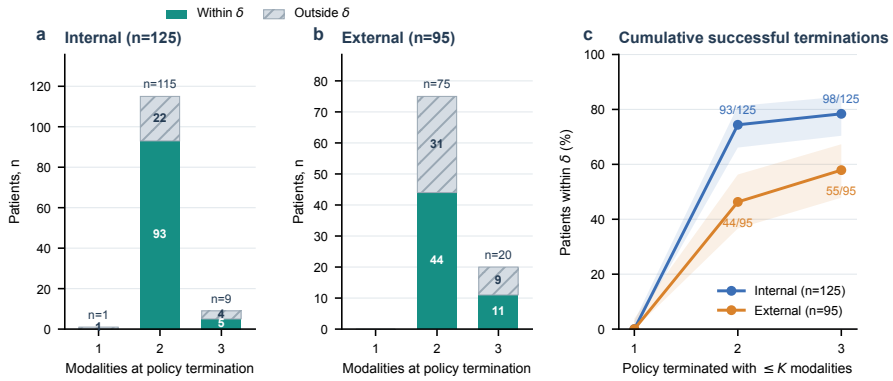

Figure 5: Patient level stopping outcomes relative to the complete four modality Dice reference, where $\delta = 0.010$ denotes the Dice tolerance used to categorize the final prediction as within or outside the complete modality reference. The figure also shows cumulative coverage by the number of acquired modalities.

prospective studies should account for acquisition time, image quality, and clinical endpoints, and should evaluate transfer to diagnosis, grading, and treatment response assessment.

7 Conclusion

We introduce world modeling from the observer's perspective, in which a model anticipates, before new evidence is acquired, the task state that would be formed by different actions that acquire information. In multisequence brain tumor MRI, NeuroBelief predicts Pathology Beliefs conditioned on the current patient state and candidate sequence and uses them to support adaptive acquisition for individual patients. The results support state prediction conditioned on the candidate as a directly testable object for adaptive acquisition, with downstream performance that varies across cohorts and operating points.

AI USE STATEMENT

Generative AI tools were used for language editing, LaTeX assistance, and consistency checks of the manuscript and reported tables. They were not used to generate experimental data or execute experiments. The authors independently verified all AI-assisted revisions against the source results and literature and take full responsibility for the paper.

ETHICS STATEMENT

We have ensured that all data used in this research comply with ethical guidelines and proper citations. The study follows responsible AI practices and avoids any data manipulation or misrepresentation.

REPRODUCIBILITY STATEMENT

The experiments in this work are fully reproducible. All datasets and code used for model training and evaluation are publicly available. Detailed information on the experimental setup, model configurations, and hyperparameters is provided to enable independent reproduction of results.

REFERENCES

M. Adewole, J. D. Rudie, A. Gbadamosi, et al. The brats-africa dataset: Expanding the brain tumor segmentation data to capture african populations. *Radiology: Artificial Intelligence*, 7(4): e240528, 2025.

BraTS 2023 Challenge. Asnr-miccai brain tumor segmentation (brats) 2023: Adult glioma segmentation. Synapse, syn51156910, 2023.

Yuhang Ding, Xin Yu, and Yi Yang. Rfnet: Region-aware fusion network for incomplete multimodal brain tumor segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3975–3984, 2021.

S. Eskreis-Winkler, A. Bhowmik, L. H. Kelly, et al. Adaptive breast mri scanning using ai. *Radiology*, 315(3):e242924, 2025.

O. B. Guney, K. S. Saichandran, K. Elzokm, Z. Zhang, and V. B. Kolachalama. Active feature acquisition via explainability-driven ranking. In *International Conference on Machine Learning*, pp. 20748–20765. PMLR, 2025.

Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pp. 2555–2565. PMLR, 2019.

Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2020.

Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for continuous control. In *International Conference on Learning Representations*, 2024.

H. Jiang, Z. Sun, N. Jia, M. Li, Y. Sun, S. Luo, S. Song, and G. Huang. Cardiac copilot: Automatic probe guidance for echocardiography with world model. In *Medical Image Computing and Computer Assisted Intervention*, pp. 190–199, 2024.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.

J. Kong, Y. He, C. Xia, R. Ge, and S. Li. Mri contrast enhancement kinetics world model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1288–1299, 2026.

Abhijit Konwer, Xiaoyang Hu, Jaemin Bae, Xuan Xu, Chao Chen, and Prateek Prasanna. Enhancing modality-agnostic representations via meta-learning for brain tumor segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 21415–21425, 2023.

S. Li, C. Chen, and J. Han. Simmlm: A simple framework for multi-modal learning with missing modality. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 24068–24077, 2025.

Y. Li and J. Oliva. Active feature acquisition with generative surrogate models. In *International Conference on Machine Learning*, pp. 6450–6459. PMLR, 2021.

Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems*, volume 14, 2001.

Hong Liu, Dong Wei, Dong Lu, Jun Sun, Li Wang, and Yefeng Zheng. M3ae: Multimodal representation learning for brain tumor segmentation with missing modalities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 1657–1665, 2023.

Chao Ma, Sebastian Tschiatschek, Konstantina Palla, Jose Miguel Hernandez-Lobato, Sebastian Nowozin, and Cheng Zhang. Eddi: Efficient dynamic discovery of high-value information with partial vae. In *International Conference on Machine Learning*, pp. 4234–4243. PMLR, 2019.

Alexander L. I. Norcliffe, C. Lee, F. Imrie, M. van der Schaar, and P. Lio. Stochastic encodings for active feature acquisition. In *International Conference on Machine Learning*, pp. 46784–46814. PMLR, 2025.

Y. Qiu, K. Jiang, H. Yao, Z. Wang, and S. Satoh. Does adding a modality really make positive impacts in incomplete multi-modal brain tumor segmentation? *IEEE Transactions on Medical Imaging*, 44(5):2194–2205, 2025.

H. Shim, S. J. Hwang, and E. Yang. Joint active feature acquisition and classification with variable-size set encoding. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Peibo Song, Xiaotian Xue, Jinshuo Zhang, Zihao Wang, Jinhua Liu, Shujun Fu, Fangxun Bao, and Si Yong Yeo. Uni-encoder meets multi-encoders: Representation before fusion for brain tumor segmentation with missing modalities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15627–15638, 2026.

Steffi C. Thust, Sabine Heiland, Andrea Falini, Hans Rolf Jäger, Adam D. Waldman, Pia C. Sundgren, Carlo Godi, Vasileios K. Katsaros, Ana Ramos, Nuria Bargallo, Meike W. Vernooij, Tarek Yousry, Martin Bendszus, and Marion Smits. Glioma imaging in europe: A survey of 220 centres and recommendations for best clinical practice. *European Radiology*, 28(8):3306–3317, 2018. doi: 10.1007/s00330-018-5314-5.

M. Valancius, M. Lennon, and J. Oliva. Acquisition conditioned oracle for nongreedy active feature acquisition. In *International Conference on Machine Learning*, pp. 48957–48975. PMLR, 2024.

Hu Wang, Yiming Chen, Chuan Ma, James Avery, Louise Hull, and Gustavo Carneiro. Multi-modal learning with missing modality via shared-specific feature modelling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15878–15887, 2023.

Michael Weller, Martin van den Bent, Matthias Preusser, et al. Eano guidelines on the diagnosis and treatment of diffuse gliomas of adulthood. *Nature Reviews Clinical Oncology*, 18:170–186, 2021. doi: 10.1038/s41571-020-00447-z.

Patrick Y. Wen, Martin van den Bent, G. Youssef, et al. Rano 2.0: Update to the response assessment in neuro-oncology criteria for high- and low-grade gliomas in adults. *Journal of Clinical Oncology*, 41(33):5187–5199, 2023. doi: 10.1200/JCO.23.01059.

H. Yang, J. Sun, and Z. Xu. Learning unified hyper-network for multi-modal mr image synthesis and tumor segmentation with missing modalities. *IEEE Transactions on Medical Imaging*, 42 (12):3678–3689, 2023.

Y. Yang, Z.-Y. Wang, Q. Liu, S. Sun, K. Wang, R. Chellappa, Z. Zhou, A. Yuille, L. Zhu, Y.-D. Zhang, and J. Chen. Medical world model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8319–8329, 2025.

Z. Yang, G. Wang, J. Hendler, M. K. Kalra, and P. Yan. X-win: Building chest radiograph world model via predictive sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6920–6930, 2026.

Yao Zhang, Nanjun He, Jiawei Yang, Yuexiang Li, Dong Wei, Yawen Huang, Yang Zhang, Zhiqiang He, and Yefeng Zheng. mmformer: Multimodal medical transformer for incomplete multimodal learning of brain tumor segmentation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pp. 107–117. Springer, 2022.

S. Zhu, Y. Chen, W. Chen, S. Jiang, G. Zhou, Y. Wang, F. Qin, C. Wang, and Q. Tian. No modality left behind: Adapting to missing modalities via knowledge distillation for brain tumor segmentation. *Medical Image Analysis*, 112:104108, 2026.

A Learning Objectives

NeuroBelief is trained in three stages. Stage I learns the Pathology Belief from incomplete MRI observations. Stage II learns the transition associated with acquiring a candidate sequence. Stage III estimates task performance from current and predicted future states for acquisition decisions. Components learned in earlier stages are frozen before training the subsequent stage.

Stage I: Pathology Belief. The Belief Encoder and Segmentation Head are trained with

$$\mathcal{L}_{\text{belief}} = \mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{Dice}} + \lambda_{\text{align}}\mathcal{L}_{\text{align}}, \tag{16}$$

where $\lambda_{\text{align}} = 0.1$. For two observation subsets S_i and S_j from the same patient, we remove the spatial mean independently from each belief channel and denote the resulting representations by $\widetilde{B}^{S_i}$ and $\widetilde{B}^{S_j}$. Their alignment loss is

$$\mathcal{L}_{\text{align}} = 1 - \cos\left(\text{vec}(\widetilde{B}^{S_i}), \text{vec}(\widetilde{B}^{S_j})\right). \tag{17}$$

This encourages representations from different observation subsets of the same patient to remain comparable in the shared representation space, while the segmentation objective preserves differences that are relevant to the task.

Stage II: Candidate transition. The Belief Encoder is frozen. The Action Encoder and Spatial Dynamics predict the Pathology Belief after candidate acquisition:

$$\mathcal{L}_{\text{dyn}} = \mathcal{L}_{\cos} + \lambda_{\text{mse}}\mathcal{L}_{\text{mse}}, \tag{18}$$

with

$$\mathcal{L}_{\cos} = 1 - \frac{1}{|\Omega|}\sum_{v\in\Omega}\left\langle \text{Norm}_C(\widehat{B}^{(a)}_{t+1,v}), \text{Norm}_C(B^{(a)}_{t+1,v})\right\rangle, \tag{19}$$

and

$$\mathcal{L}_{\text{mse}} = \text{MSE}\left(\widehat{B}^{(a)}_{t+1}, B^{(a)}_{t+1}\right). \tag{20}$$

Stage III: Acquisition value. The Belief Encoder, Segmentation Head, Action Encoder, and Spatial Dynamics are frozen. Only the Value Head is optimized. The ground truth segmentation is used to construct supervised Dice targets during training, but it is not available to the Value Head or the acquisition policy during inference.

For predicted and target Dice vectors $\hat{\mathbf{d}}$ and $\mathbf{d}$, we use

$$\ell_{\text{reg}}(\hat{\mathbf{d}}, \mathbf{d}) = \text{SmoothL1}(\hat{\mathbf{d}}, \mathbf{d}) + \lambda_{\text{mean}}\text{SmoothL1}\left(\text{Mean}(\hat{\mathbf{d}}), \text{Mean}(\mathbf{d})\right). \tag{21}$$

The current, future, and gain losses are

$$\mathcal{L}_{\text{cur}} = \ell_{\text{reg}}(\hat{\mathbf{d}}_t, \mathbf{d}_t), \tag{22}$$

$$\mathcal{L}_{\text{future}} = \ell_{\text{reg}}(\hat{\mathbf{d}}_{t+1}^{(a)}, \mathbf{d}_{t+1}^{(a)}), \tag{23}$$

$$\mathcal{L}_{\text{gain}} = \ell_{\text{reg}}\left(\hat{\mathbf{d}}_{t+1}^{(a)} - \hat{\mathbf{d}}_t, \mathbf{d}_{t+1}^{(a)} - \mathbf{d}_t\right). \tag{24}$$

For candidate a, let

$$q_a = \text{Mean}(\mathbf{d}_{t+1}^{(a)}), \qquad \hat{q}_a = \text{Mean}(\hat{\mathbf{d}}_{t+1}^{(a)}). \tag{25}$$

Relative candidate values are supervised by

$$\mathcal{L}_{\text{adv}} = \text{SmoothL1}\left(\hat{\mathbf{q}} - \text{Mean}(\hat{\mathbf{q}}), \mathbf{q} - \text{Mean}(\mathbf{q})\right). \tag{26}$$

For ranking, we use pairs satisfying

$$\mathcal{P} = \{(i,j) : |q_i - q_j| > \epsilon_{\text{rank}}\}, \qquad s_{ij} = \text{sign}(q_i - q_j), \tag{27}$$

and

$$\mathcal{L}_{\text{rank}} = \frac{1}{|\mathcal{P}|} \sum_{(i,j)\in\mathcal{P}} \text{softplus}\left(-\frac{s_{ij}(\hat{q}_i - \hat{q}_j)}{T}\right). \tag{28}$$

When $\mathcal{P}$ is empty, $\mathcal{L}_{\text{rank}} = 0$.

The complete objective is

$$\mathcal{L}_{\text{value}} = \mathcal{L}_{\text{cur}} + \mathcal{L}_{\text{future}} + \mathcal{L}_{\text{gain}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{rank}}. \tag{29}$$

All loss terms in the Value Head objective are assigned a weight of 1.

B Experimental Implementation

The four MRI sequences are T1n, T1c, T2w, and T2 FLAIR. The input contains four channels in a fixed order together with the presence mask m_t. Missing sequences are replaced with zeros. Stage I covers all $2^4 - 1 = 15$ nonempty observation subsets.

For transition learning, each noncomplete subset defines one transition for every remaining candidate. States with one, two, and three observed sequences therefore provide

$$4 \times 3 + 6 \times 2 + 4 \times 1 = 28 \tag{30}$$

legal transitions per patient.

The Belief Encoder and Segmentation Head are trained for 60 epochs. Spatial Dynamics and the Action Encoder are then trained for 20 epochs with the Belief Encoder frozen. The Value Head is trained for 20 epochs with all preceding components frozen. We use AdamW with learning rates 2×10^{-4} for Stage I and 1×10^{-4} for Stages II and III, and weight decay 10^{-5}. Checkpoints and acquisition settings are selected using the development cohort. The internal test and external cohorts are used only for evaluation.

C Candidate State Prediction

Evaluation metrics. The current, predicted, and realized future beliefs are decoded with the same frozen Segmentation Head:

$$P_t = H_\eta(B_t), \qquad \widehat{P}_{t+1}^{(a)} = H_\eta(\widehat{B}_{t+1}^{(a)}), \qquad P_{t+1}^{(a)} = H_\eta(B_{t+1}^{(a)}). \tag{31}$$

Their signed probability changes are

$$\Delta P^{(a)} = P_{t+1}^{(a)} - P_t, \qquad \Delta\widehat{P}^{(a)} = \widehat{P}_{t+1}^{(a)} - P_t. \tag{32}$$

For $C = 3$ output channels and spatial domain Ω,

$$\text{ChangeMAE}(a) = \frac{1}{C|\Omega|} \sum_{c,v} \left|\Delta\widehat{P}_{c,v}^{(a)} - \Delta P_{c,v}^{(a)}\right|. \tag{33}$$

Spatial Pearson is computed between the absolute change maps

$$M^{(a)}(v) = \frac{1}{C}\sum_{c} |\Delta P^{(a)}_{c,v}|, \qquad \widehat{M}^{(a)}(v) = \frac{1}{C}\sum_{c} |\Delta \widehat{P}^{(a)}_{c,v}|. \tag{34}$$

Top 10% Dice measures the overlap between locations with the largest 10% of values in these maps.

For retrospective task consequence evaluation, let $\mathbf{d}_t$, $\widetilde{\mathbf{d}}^{(a)}_{t+1}$, and $\mathbf{d}^{(a)}_{t+1}$ denote WT, TC, and ET Dice obtained from P_t, $\widehat{P}^{(a)}_{t+1}$, and $P^{(a)}_{t+1}$, respectively. All are computed against the ground truth segmentation. We define

$$\widetilde{g}_t(a) = \mathrm{Mean}\left(\widetilde{\mathbf{d}}^{(a)}_{t+1} - \mathbf{d}_t\right), \qquad g_t(a) = \mathrm{Mean}\left(\mathbf{d}^{(a)}_{t+1} - \mathbf{d}_t\right). \tag{35}$$

Gain MAE is

$$\mathrm{GainMAE} = \frac{1}{|\mathcal{A}_t|}\sum_{a\in\mathcal{A}_t} |\widetilde{g}_t(a) - g_t(a)|. \tag{36}$$

We further report

$$a^{\mathrm{pred}} = \arg\max_a \widetilde{g}_t(a), \qquad a^{\mathrm{oracle}} = \arg\max_a g_t(a), \tag{37}$$

with

$$\mathrm{Top1} = \mathbb{I}[a^{\mathrm{pred}} = a^{\mathrm{oracle}}], \qquad \mathrm{Regret} = g_t(a^{\mathrm{oracle}}) - g_t(a^{\mathrm{pred}}). \tag{38}$$

The ground truth segmentation is used only to construct supervised targets and retrospective evaluation metrics. For cohort statistics, matched transitions are first averaged within each patient, after which the mean and sample standard deviation are computed across patients.

Controls. Identity assumes

$$\widehat{B}^{(a)}_{t+1} = B_t. \tag{39}$$

Spatial Population uses the mean training update for each observation subset S and candidate a:

$$\overline{\Delta B}_{S,a} = \frac{1}{|\mathcal{D}_{\mathrm{train}}|}\sum_{i\in\mathcal{D}_{\mathrm{train}}} \left[B_i(S\cup\{a\}) - B_i(S)\right], \tag{40}$$

and predicts

$$\widehat{B}_i(S,a) = B_i(S) + \overline{\Delta B}_{S,a}. \tag{41}$$

Global removes the spatial layout of the current belief while preserving channel information and candidate conditioning:

$$G(B_t) = \mathrm{Broadcast}\left(\frac{1}{D_B H_B W_B}\sum_{d,h,w} B_{t,:,d,h,w}\right). \tag{42}$$

It otherwise uses the same Action Encoder, Spatial Dynamics architecture, transitions, objective, optimization protocol, and spatial target as NeuroBelief.

Direct Logit Delta predicts the WT, TC, and ET logit changes directly instead of constructing a future Pathology Belief. Its target is

$$\Delta L^{(a)}_t = \mathrm{logit}\left(\mathrm{clip}(P^{(a)}_{t+1}, \epsilon_{\mathrm{clip}}, 1-\epsilon_{\mathrm{clip}})\right) - \mathrm{logit}\left(\mathrm{clip}(P_t, \epsilon_{\mathrm{clip}}, 1-\epsilon_{\mathrm{clip}})\right), \tag{43}$$

where $\epsilon_{\mathrm{clip}} = 10^{-6}$. It is trained with

$$\mathcal{L}_{\mathrm{logit}} = \mathrm{SmoothL1}_{\beta=0.1}(\widehat{\Delta L}, \Delta L) + 0.1\left[1 - \cos(\widehat{\Delta L}, \Delta L)\right]. \tag{44}$$

At evaluation time,

$$\widehat{P}^{(a)}_{t+1,\mathrm{direct}} = \sigma\left(\mathrm{logit}\left(\mathrm{clip}(P_t, \epsilon_{\mathrm{clip}}, 1-\epsilon_{\mathrm{clip}})\right) + \widehat{\Delta L}^{(a)}_t\right). \tag{45}$$

This reconstructed output is evaluated with the same probability change and task consequence metrics as NeuroBelief.

Table 4: Additional controls for candidate state prediction. Values are reported as the mean ± sample standard deviation across patients.

A. Global spatial control			
Cohort	**Predictor**	**Change MAE ($\times 10^{-3}$) ↓**	**Pearson ↑**
BraTS 2023	Global	24.199 ± 13.578	0.409 ± 0.088
BraTS Africa	Global	40.275 ± 21.711	0.440 ± 0.067
B. Latent candidate identification			
Cohort	**Predictor**	**Centered L1 ↑**	**Centered cosine ↑**
BraTS 2023	Chance	0.417 ± 0.000	0.417 ± 0.000
	Global	0.484 ± 0.062	0.423 ± 0.017
	Spatial Population	0.466 ± 0.047	0.761 ± 0.090
	NeuroBelief	$\mathbf{0.574 \pm 0.082}$	$\mathbf{0.818 \pm 0.095}$
BraTS Africa	Chance	0.417 ± 0.000	0.417 ± 0.000
	Global	0.458 ± 0.051	0.424 ± 0.019
	Spatial Population	0.462 ± 0.045	0.713 ± 0.086
	NeuroBelief	$\mathbf{0.518 \pm 0.073}$	$\mathbf{0.772 \pm 0.099}$

Conditioning tests. Patient permutation keeps the observation subset and candidate fixed while replacing the current belief with that of another patient. Candidate permutation keeps the current belief fixed while cyclically replacing the candidate identity. In both cases, predictions are evaluated against the consequence associated with the original patient and candidate.

Latent candidate identification. We evaluate states with $K = 1$ or $K = 2$ observed sequences. States with $K = 3$ are excluded because only one candidate remains. For candidate a, the realized and predicted belief changes are

$$\Delta B^{(a)} = B^{(a)}_{t+1} - B_t, \qquad \Delta \widehat{B}^{(a)} = \widehat{B}^{(a)}_{t+1} - B_t. \tag{46}$$

After removing the spatial mean independently from each belief channel, candidate identity is recovered by matching each predicted change to the realized candidate changes using centered L1 distance or centered cosine similarity. Identification is correct when the matched candidate equals the candidate used for prediction. The chance reference is $1/3$ for $K = 1$ and $1/2$ for $K = 2$, giving an average reference of 0.417.

Global loses substantial spatial predictive accuracy and remains near chance under centered cosine identification. Spatial Population retains population updates associated with each candidate, while NeuroBelief further improves latent candidate identification on both cohorts.

Relative to Spatial Population, NeuroBelief reduces paired Change MAE by 0.476×10^{-3} on BraTS 2023 (95% CI, 0.409–0.544×10^{-3}) and 0.589×10^{-3} on BraTS Africa (0.457–0.722×10^{-3}). Relative to Direct Logit Delta, the paired differences in Gain MAE, Top 1 accuracy, and regret have 95% confidence intervals excluding zero on both cohorts.

D Generalization to Withheld Transition Combinations

Protocol. A transition type is defined by the current observation set S and an unobserved candidate modality a:

$$\mathcal{T}_{1,2} = \{(S, a) : |S| \in \{1, 2\},\ a \notin S\}. \tag{47}$$

The four modalities produce 24 transition types for $K = 1$ and $K = 2$. We divide them into three disjoint folds. Each fold removes eight transition types from Dynamics training: four with $K = 1$, four with $K = 2$, and two acquisitions of each candidate modality. The remaining 16 types and

Table 5: Prediction on transition combinations excluded from Dynamics training. Values are reported as the mean ± sample standard deviation across patients.

Cohort	Method	Latent L1 ↓	Change MAE ($\times 10^{-3}$) ↓	Direction ↑	Pearson ↑	Top 10% Dice ↑	Decoded Candidate ID ↑
BraTS 2023	Seen Spatial Population	0.159 ± 0.003	5.228 ± 2.948	0.142 ± 0.069	0.552 ± 0.039	0.855 ± 0.048	0.520 ± 0.095
	NeuroBelief	$\mathbf{0.111 \pm 0.004}$	$\mathbf{4.932 \pm 2.995}$	$\mathbf{0.233 \pm 0.063}$	$\mathbf{0.568 \pm 0.035}$	$\mathbf{0.874 \pm 0.041}$	$\mathbf{0.576 \pm 0.085}$
BraTS Africa	Seen Spatial Population	0.164 ± 0.006	11.673 ± 7.251	0.110 ± 0.068	0.543 ± 0.037	0.838 ± 0.063	0.498 ± 0.089
	NeuroBelief	$\mathbf{0.118 \pm 0.007}$	$\mathbf{11.477 \pm 7.271}$	$\mathbf{0.177 \pm 0.083}$	$\mathbf{0.551 \pm 0.035}$	$\mathbf{0.847 \pm 0.063}$	$\mathbf{0.541 \pm 0.092}$

all four $K = 3$ transitions remain available during training. Model selection uses only transitions available during training.

Seen Spatial Population is estimated from available training transitions with the same observation size K and candidate modality a and never uses the removed transition type.

Latent L1 is

$$\mathrm{LatentL1}(a) = \frac{1}{C_B|\Omega_B|} \left\| \widehat{B}^{(a)}_{t+1} - B^{(a)}_{t+1} \right\|_1 . \tag{48}$$

Change MAE, Spatial Pearson, and Top 10% Dice follow Appendix C. Direction measures elementwise agreement in the sign of decoded probability change:

$$\mathrm{Direction}(a) = \frac{1}{C|\Omega|} \sum_{c,v} \mathbb{I}\left[\mathrm{sign}(\Delta \widehat{P}^{(a)}_{c,v}) = \mathrm{sign}(\Delta P^{(a)}_{c,v}) \right] . \tag{49}$$

Decoded Candidate ID matches the complete predicted signed probability change to the realized candidate changes using cosine distance. For all reported values, transitions and folds are first averaged within each patient before cohort statistics are computed.

NeuroBelief improves over Seen Spatial Population on both cohorts. The largest differences occur in latent prediction error and direction agreement, while decoded spatial metrics and Candidate ID show the same trend. These results test generalization to (S, a) transition configurations excluded from Dynamics optimization.

E Patient Level Closed Loop Adaptive Acquisition

Experimental setting. The closed loop starts from T1c, selected on the development cohort. At each step, NeuroBelief evaluates the remaining candidates and either acquires one sequence or stops. After acquisition, the selected image is added to the observation set and the Pathology Belief is recomputed. Final segmentation is decoded from the acquired sequences.

The primary stopping threshold is 0.03, selected on the development cohort before test evaluation. Population Fixed applies the same acquisition plan to every patient. Spatial Population uses state changes estimated from the training cohort and does not use transitions specific to individual patients. Input containing all four modalities is used only as a performance reference.

Oracle opportunity. We use a retrospective Oracle to quantify the opportunity available when acquisition is tailored to individual patients under the frozen segmentation model. For each patient, the Oracle selects the acquisition path with the highest realized Dice. The retrospective best fixed strategy is the single fixed path with the highest cohort mean realized Dice under the same maximum budget. Both references use realized test outcomes only for this analysis and are not deployable acquisition policies.

With early stopping, the paired Dice difference is 0.0069 (95% CI, 0.0045–0.0099). This gap quantifies the retrospective headroom in task performance associated with selecting acquisition paths for individual patients rather than using the best single fixed path.

Matched acquisition budgets. Because the compared policies do not always use the same average number of sequences, we perform a sensitivity analysis with matched acquisition budgets. Population references are obtained by linear interpolation along their Pareto curves relating performance to acquisition cost at the same average acquisition budget. The interpolation weights are applied

Table 6: Retrospective Oracle opportunity on the BraTS 2023 internal test cohort. Values are reported as the mean ± sample standard deviation across patients.

Maximum budget	Patient Oracle	Retrospective Best Fixed	Difference
2	$\mathbf{0.8543 \pm 0.1140}$	0.8478 ± 0.1203	$+0.0064 \pm 0.0185$
3	$\mathbf{0.8601 \pm 0.1105}$	0.8521 ± 0.1176	$+0.0080 \pm 0.0182$
4 with early stopping	$\mathbf{0.8611 \pm 0.1098}$	0.8541 ± 0.1170	$+0.0069 \pm 0.0158$

Table 7: Budget matching sensitivity analysis at representative operating points. Values are reported as the mean ± sample standard deviation across patients.

Cohort	Threshold	Sequence reduction	NeuroBelief Dice	Δ vs. Population Fixed	Δ vs. Spatial Population
BraTS 2023	0.030	$48.40 \pm 6.92\%$	0.8473 ± 0.1294	-0.0007 ± 0.0343	-0.0026 ± 0.0361
	0.010	$45.20 \pm 10.38\%$	0.8501 ± 0.1170	$+0.0016 \pm 0.0138$	-0.0004 ± 0.0090
	0.000	$39.60 \pm 13.54\%$	0.8502 ± 0.1175	$+0.0007 \pm 0.0114$	-0.0007 ± 0.0073
BraTS Africa	0.030	$44.74 \pm 10.25\%$	0.8192 ± 0.1276	$+0.0070 \pm 0.0408$	$+0.0035 \pm 0.0290$
	0.010	$39.21 \pm 12.45\%$	0.8229 ± 0.1266	$+0.0069 \pm 0.0319$	$+0.0071 \pm 0.0292$
	0.000	$30.53 \pm 10.43\%$	0.8252 ± 0.1266	$+0.0032 \pm 0.0179$	$+0.0044 \pm 0.0232$

to Dice values for individual patients before paired differences are computed. These references are used only for sensitivity analysis and do not represent executable acquisition policies.

The matched budget differences remain small on BraTS 2023 and positive at the representative operating points on BraTS Africa. The primary comparison is therefore not explained solely by differences in average acquisition budget.

Table 8: Acquisition paths after the common T1c starting point at threshold 0.03.

Added path	BraTS 2023, n (%)	BraTS Africa, n (%)
T2 FLAIR	114 (91.2)	73 (76.8)
T2 FLAIR → T2w	6 (4.8)	19 (20.0)
T2w → T2 FLAIR	3 (2.4)	1 (1.1)
T2w	1 (0.8)	2 (2.1)
STOP	1 (0.8)	0 (0.0)

Acquisition paths for individual patients. Table 8 reports the paths selected after the common T1c starting point at the primary threshold. Most patients acquire T2 FLAIR first after T1c. The main variation is whether T2w is subsequently acquired, which occurs more frequently in the external cohort.