

# Towards user-centered interactive medical image segmentation in VR with an assistive AI agent

Pascal Spiegler, Arash Harirpoush, and Yiming Xiao

**Abstract**—Crucial in disease analysis and surgical planning, manual segmentation of volumetric medical scans (e.g. MRI, CT) is laborious, error-prone, and challenging to master, while fully automatic algorithms can benefit from user feedback. Therefore, with the complementary power of the latest radiological AI foundation models and virtual reality (VR)’s intuitive data interaction, we propose SAMIRA, a novel conversational AI agent that assists users with localizing, segmenting, and visualizing 3D medical concepts in VR. Through speech-based interaction, the agent helps users understand radiological features, locate clinical targets, and generate segmentation masks that can be refined with just a few point prompts. The system also supports true-to-scale 3D visualization of segmented pathology to enhance patient-specific anatomical understanding. Furthermore, to determine the optimal interaction paradigm under near-far attention-switching for refining segmentation masks in an immersive, human-in-the-loop workflow, we compare VR controller pointing, head pointing, and eye tracking as input modes. With a user study, evaluations demonstrated a high usability score (SUS=90.0±9.0), low overall task load, as well as strong support for the proposed VR system’s guidance, training potential, and integration of AI in radiological segmentation tasks.

**Index Terms**—Medical image segmentation, Virtual reality, Human-in-the-loop, AI agent, Foundation model, Attention switching, Eye tracking, Medical visualization, Clinical decision support

## I. INTRODUCTION

Medical image segmentation is a critical task in clinical diagnosis and treatment planning, particularly for identifying and quantifying abnormalities such as tumors, stroke lesions, and other pathological anomalies. The process involves producing segmentation masks that delineate and highlight regions of interest to provide a basis for further analysis, treatment decisions, and longitudinal tracking. However, traditional workflows for medical image segmentation are time-consuming and labor-intensive, typically requiring experts to manually annotate up to hundreds of 2D slices to isolate structures within 3D MRI or CT scans. Furthermore, while segmenting pathologies (e.g., tumor) accurately is itself a demanding skill, it often requires extensive hours of supervision and training to develop diagnostic confidence and anatomical precision [1]. Finally, visualizations of these annotations are similarly challenging: clinicians either scroll through superimposed binary masks on 2D slices across the axial, sagittal, and coronal planes, or view

3D reconstructions rendered on flat screens, both of which lack real spatial context and a true sense of scale.

While virtual reality (VR) can offer more intuitive 3D medical data visualization and interaction, especially under high spatial constraints (e.g., in the clinic), recent developments in foundation artificial intelligence (AI) models, such as vision-language models (VLMs) have demonstrated early promise to further enhance the efficiency, accuracy, and interactivity for tasks in VR in the form of AI agents [2], [3]. As an alternative to conventional brush painting-based segmentation paradigms, imagine a workflow that unites the efficiency of AI with the spatial interaction advantages of a virtual environment: a user reviews a brain tumor MRI in VR assisted by an AI agent. First, the agent guides them towards a representative tumor slice. When the user confirms this slice contains a tumor, they issue a simple voice command, triggering the AI agent to highlight the tumor and provide a detailed patient-specific radiological description. Crucially, this interaction goes beyond the role of a typical conversational assistant: rather than passively responding to dialogue, the AI agent actively executes a sequence of actions, from volumetric segmentation to case-based guidance and full 3D rendering with minimal input. If necessary, the user can easily refine the segmentation using natural inputs, including head pointing, gaze, or handheld controllers, and finally, the corrected mask is rendered in 3D at true scale, offering enhanced spatial interpretation of the tumor’s dimensions. Notably, as current deep learning algorithms for radiological segmentation still require human quality assurance [4]–[6], such a human-in-the-loop approach preserves user oversight while dramatically improving speed, accuracy, and spatial understanding through immersive AI assistance.

In our work, we present such a VR-based, AI-assisted medical image segmentation system that supports medical image review for both diagnostic decision-making and education. To the best of our knowledge, this system is the first of its kind, not only automating segmentation across image slices via a conversational AI agent, but also investigating optimal interaction paradigms for human-in-the-loop mask refinement. Our novel contributions are as follows: **First**, we systematically investigate optimal interaction paradigms for human-in-the-loop segmentation refinement that transitions between proximal corrections and distal menu interactions, comparing natural inputs from handheld controllers, head pointing, and gaze. **Second**, we introduce a new 3D segmentation algorithm based on the BiomedParse and SAM2 foundation models which reduces mask drift from noise during slice-to-slice annotation propagation in medical images. **Finally**, we present

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

Pascal Spiegler, Arash Harirpoush, and Yiming Xiao are with the Department of Computer Science and Software Engineering, Concordia University, Montreal, Canada. E-mails: pascal.spiegler@mail.concordia.ca, aharirpoush@gmail.com, yiming.xiao@concordia.ca

SAMIRA, a conversational AI agent that assists with segmenting 3D radiological scans via voice commands, provides interactive guidance, supports iterative refinement to preserve expert oversight, and enables life-scale 3D visualization of results in immersive VR.

## II. RELATED WORK

Although research on AI assistants for clinical training in VR is still in its infancy, early work has begun to demonstrate its potential to support clinical workflows and healthcare education. For example, Liaw et al. [7] developed a conversational AI assistant within a VR simulation for sepsis training, which matched human-controlled scenarios in clinical and communication performance and led to significantly higher test scores. Furthermore, Chheang et al. [8] introduced a generative AI-based virtual assistant in a VR anatomy education environment, highlighting the potential of such assistants to support anatomy education. Although these systems leverage large language models (LLMs) to provide educational support via dialogue, they primarily function as virtual assistants focused on delivering information. In contrast, AI agents should autonomously execute sequential tasks in response to high-level user prompts, for example, as in our scenario, performing segmentation, providing clinical context via text and images, and rendering 3D models. Retrieval-Augmented Generation (RAG) often powers such agents by retrieving relevant examples from a structured knowledge database and combining similarity search results with generative models to offer refined responses to queries. This contextualized guidance is particularly useful for medical tasks, where visual patterns vary across cases. To our knowledge, no prior work has integrated AI agents into immersive VR for interactive 3D radiological segmentation.

Recent development in foundation models has paved the way for interactive medical image analysis with multi-modal inputs. One such model, BiomedParse [9], is a Transformer-based vision-language model that uses separate encoders for medical images and clinical text data. Trained on 1.2 million paired 2D radiological images and reports, it supports detection, classification, and segmentation of 82 clinical concepts (e.g. tumor) across 9 imaging modalities using natural language prompts. On the other hand, SAM2 [10] is a foundation model for segmenting objects in images and video using user-provided point and/or box prompts that enables prompt-based segmentation mask refinement and leverages a memory encoder to propagate masks across video frames. Although developed for natural image domains, SAM2’s ability to handle sequential image data positions itself as a promising tool for 3D medical imaging segmentation, where volumetric scans, such as CT or MRI can be viewed as series of 2D images analogous to videos [11], [12]. This opens up the possibility of combining BiomedParse’s language-driven label mask generation with SAM2’s interactive label refinement and propagation capabilities to enable human-in-the-loop segmentation workflows for 3D medical imaging, an underexplored area, particularly in immersive VR environments. While prior VR segmentation systems have allowed users to paint regions

of interest using hand gestures [13], [14] or have used predefined anatomy-specific deep learning models [15], none have integrated foundation segmentation models with agent-driven guidance to support interactive refinement.

In immersive medical image segmentation with SAM2, users will need to frequently switch their visual and motor attention between the image being annotated and a spatially decoupled user interface (UI) menu for triggering actions (e.g., toggling between negative/positive prompts, resetting prompts, etc.), thus creating a dual-focus challenge. Previously, Rashid et al. [16] compared proximal (on-device) versus distal (remote) widget selection in distributed user interfaces and found that proximal methods were faster and preferred for complex, multi-step tasks, while distal methods yielded lower error rates for simpler interactions, suggesting that high-precision tasks like segmentation prompt placement may benefit from proximal selections and that simpler menu selections may benefit from being performed at a distance. In our case, the interactive refinement of AI-predicted segmentations with SAM2 in VR necessitates identifying optimal interaction paradigms for completing the task under attention-switching. *However, the ideal interaction paradigm (e.g., controller, head pointing, or eye-tracking) for attention switching between proximal and distal displays remains unexplored.* Sidenmark et al. [17] conducted a head-mounted VR study comparing gaze, head, and controller pointing for dynamically revealed target selection, showing that both gaze- and controller-based pointing significantly outperformed head pointing in terms of speed and precision, though they did not explore static dual-panel switching scenarios. Luro et al. [18] compared eye tracking with hand-controller aiming tasks in VR and reported that controllers achieved the highest placement accuracy, while gaze-based selection felt more natural, but suffered from increased selection accuracy variability. Studies of multi-depth targeting, such as the experiment of Schultheis et al. [19] revealed that eye-based selection can achieve higher throughput across varying depth planes, but controller input offers more stable performance when precision is critical. Furthermore, Xu et al.’s evaluation of text-selection techniques in VR [20] demonstrated that head-pointing with click confirmation strikes a balance between speed and accuracy, ranking just behind controller pointing in speed while maintaining a minimal task load, positioning it as a possibly viable paradigm for attention switching interactions. With these previous insights, we will investigate the optimal interaction paradigms for point placement and menu selection under attention switching conditions for our application.

## III. SYSTEM OVERVIEW

To address the aforementioned issues, we present SAMIRA (Segmentation Assistant for Multimodal Interaction and Radiological Analysis), a novel conversational agent designed to support human-in-the-loop 3D medical image segmentation by generating segmentation masks, enabling efficient mask refinement, and providing radiological guidance through speech and reference images. Figure 1 illustrates the system’s key components during the segmentation of a liver tumor in a CT scan.

All VR development was conducted in Unity 2022.3.f1 on a desktop equipped with an NVIDIA RTX 3090 and an HTC Vive Pro Eye VR headset. In the VR system, medical image slices are displayed on a virtual panel anchored to a VR controller in the user’s hand, allowing flexible repositioning for detailed inspection. Users navigate through image slices by rotating the thumbpad on the HTC Vive controller, clockwise to advance and counter-clockwise to reverse. Each 60° rotation corresponds to a single slice, reinforced by a haptic pulse that provides tactile feedback at each transition.

The menu interface is structured into three panels (see Figure 1): on the left panel, AI-driven guidance is provided as text and synthesized speech for the given segmentation task. In the middle panel, interactive controls with functional buttons allow users to issue voice commands, refine predicted masks, propagate segmentations across slices, and render results in 3D. On the right panel, reference images are retrieved in real time based on the user’s currently viewed slice based on RAG from an existing knowledge repository. On this panel, the system presents anatomically similar image examples with and without the target structure (e.g., tumor), helping users locate pathology and build a clearer understanding of its visual characteristics.

Using the AI agent’s guidance and slice-scrolling mechanism, users can find the target pathology, initiate segmentation via voice commands, refine masks using efficient natural inputs, then render a final segmentation into a true-to-scale 3D visualization to gain spatial understanding of pathologies.

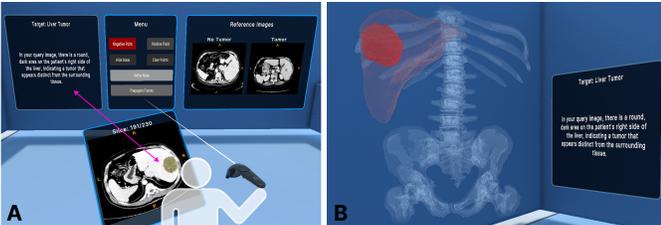


Fig. 1. A. The AI agent generates an initial segmentation of a liver tumor in CT and provides guidance using reference images and patient-specific pathology explanations. B. The final refined 3D visualization is rendered as a large, spherical, high-contrast liver tumor in red, overlaid on anatomical structures, to scale. Few refinements are expected due to the simple shape.

All real-time inference (e.g., RAG retrieval), segmentation generation, point-prompt-based segmentation refinement, mask propagation, and 3D mesh creation are realized by SAMIRA and coordinated via two dedicated Python Web-Socket servers: a local rendering server and a dedicated inference server. This is to ensure responsive, low-latency AI assistance and integration between Unity and the supporting AI models. However, future similar systems can benefit from cloud-based setups.

**Local Rendering Server:** The server ran on the same desktop as the Unity client. It receives binary segmentation masks, scales the lesions using medical image metadata, and returns a `.obj` mesh for VR visualization.

**Dedicated Inference Server:** A separate Ubuntu 22.04.2 LTS server was set up with a dedicated NVIDIA RTX 3090 to host SAMIRA’s RAG and foundational segmentation models

(BiomedParse, SAM2). It handles voice commands, reference retrieval, textual guidance, text-based segmentation generation, mask refinement, and mask propagation across frames (with SAM2).

## IV. METHODS AND MATERIALS

### A. Interaction Paradigm Evaluation

Before evaluating the full VR system with our AI agent, SAMIRA, we aimed to first reveal the optimal interaction paradigm for point prompt placement and menu selection under attention switching between the proximal handheld image display and the distal menu interface. To accomplish this, we designed an experiment to perform SAM2-based segmentation mask refinement using three distinct paradigms: *Controller*, *Head Pointing*, and *Eye Tracking*. The detailed setup is shown in Fig. 2. As a potential middle ground between *Controller* and *Eye tracking*, we hypothesize that *Head Pointing* is the optimal interaction paradigm for our intended application, with the consideration of task load, accuracy, and efficiency.

**Controller** pointing is a staple method for object and menu selection in many VR applications. We employed standard controller-based ray-casting, where a visible ray extends from the tip of the controller to intercept with the image under analysis and the menu for point prompt annotation and button clicking, respectively. Here, confirmation of selection is achieved by pressing the trigger button of the controller.

**Head Pointing** adopts an invisible ray forward from the center of the user’s headset, aligning with the head’s orientation. The interception of the ray and the image/menu is represented by a visible dot, and the controller trigger button is used for placement/selection confirmation.

**Eye Tracking** utilizes the integrated eye-tracking hardware of the HTC Vive Pro Eye headset to cast an invisible ray based on the user’s gaze direction. Similar to head pointing, a visible dot is used to indicate the gaze point on the image under analysis and the menu. Based on previous studies [21], [22], we continue to use the trigger button to confirm selection. To mitigate the adverse impacts of natural eye jitter on selection precision and user experience, we apply exponential smoothing to the normalized 3D gaze vectors during every frame using a smoothing factor of  $\alpha = 0.2$ . The gaze ray is computed as follows:

$$\mathbf{r}_{\text{smooth}} = (1 - \alpha) \cdot \mathbf{r}_{\text{previous}} + \alpha \cdot \mathbf{r}_{\text{current}} \quad (1)$$

where  $\alpha = 0.2$  is the smoothing factor,  $\mathbf{r}_{\text{previous}}$  is the smoothed direction from the previous frame, and  $\mathbf{r}_{\text{current}}$  is the normalized average of the left and right eye gaze direction vectors, originating at the midpoint of the two eyes.

The overall interaction paradigm evaluation workflow is summarized in Figure 2F, where the user begins at the middle slice of a segmented 3D scan and scrolls through the volume to identify and correct segmentation errors. Three fixed large panels (Figure 2C) are positioned three meters in front of the user, different from the panel displays for the full interactive segmentation workflow (Figure 1): the left panel displays the

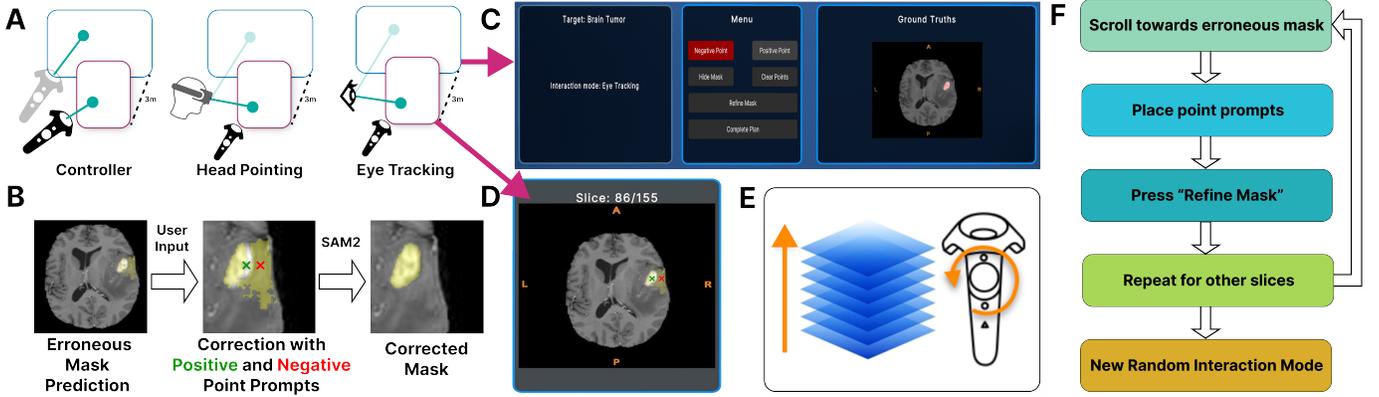


Fig. 2. System setup for user interaction paradigm evaluation under attention switching. A. Three interaction paradigms: controller ray, head pointing, and eye tracking. B. Users correct erroneous masks using positive (green) and negative (red) point prompts, refined by SAM2. C. In-VR interface for prompt selection and real-time ground truth reference. D. Medical image display with current slice and segmentation overlay. E. Controller-based slice scrolling. F. Interaction paradigm evaluation segmentation workflow.

current interaction mode (controller, head pointing, or eye tracking), the central panel provides interactive controls (i.e., functional buttons), and the right panel shows the ground truth segmentation (red, 40% opacity overlaid on the image) for the current slice under study. Here, the ground truth segmentation is used to define a consistent reference for refinement and objective evaluation of segmentation accuracy across participants in the user study. Users were instructed to correct the provided masks until they visually matched the ground truths. The central panel includes six buttons. “Positive Point” and “Negative Point” buttons allow users to place point prompts that add missing regions or remove excess segmentation, respectively. The “Hide Mask” button toggles the visibility of both the segmentation under work and ground truth to better assess tissue boundaries. The “Clear Points” button erases all placed point prompts without modifying the current mask. The “Refine Mask” button sends the current slice and point prompts to the inference server to update the segmentation. Finally, “Complete Plan” finalizes segmentation refinement and advances to the next interaction paradigm (*Controller*, *Head Pointing*, or *Eye Tracking*), selected randomly, allowing within-participant comparisons of performance across paradigms.

To evaluate interaction paradigms under attention-switching for AI-facilitated segmentation refinement, we used a T1c MRI scan from the publicly available BRATS brain tumor dataset [23], which was intensity-normalized and converted into a sequence of 155 axial JPEG image slices. All slices along with 44 slice-wise ground truth masks and 16 intentionally corrupted slice-wise masks were included. The corrupted masks represented common failure cases, such as incomplete tumor coverage, over-segmentation, missing regions, or false positives. The quality of the initial segmentation, measured using the 3D Dice score, was assessed at 0.91. 3D Dice is a standard metric for volumetric medical segmentation that quantifies the spatial overlap between the predicted mask  $A$  and the ground truth  $B$  as:

$$\text{Dice}(A, B) = \frac{2|A \cap B|}{|A| + |B|} \quad (2)$$

A Dice score of 1 indicates perfect agreement; 0 indicates no overlap.

### B. SAMIRA - User Interface and Workflow

The full workflow of our interactive medical image segmentation leverages SAMIRA, our assistive AI agent via the VR user interface described in Section III. The interface enables SAMIRA to operate on unseen cases without ground truth segmentations, and supplies reference images and case-relevant descriptive guidance in real time. This supports the potential for future clinical and educational applications beyond the controlled evaluation settings showcased in our study.

For demonstration and evaluation purposes, the segmentation workflow was applied to a brain tumor MRI scan from another dataset, Pretreat-MetsToBrain-Masks [24], and a liver tumor CT scan from the LiTS dataset [25]. These two datasets were deliberately chosen to assess generalization beyond the training distribution of the underlying models, since the model we are introducing, BiomedParse was trained on the BRATS dataset used in the interaction paradigm test. While the examples focused on tumors, the workflow is compatible with any of the 82 clinical targets supported by BiomedParse. The brain tumor example was selected as the more difficult case, with branching regions (see Figure 3E), whereas the liver tumor case was expected to be easier to segment, with a smoother, spherical structure in which less refinement is expected (see Figure 1).

The full workflow is illustrated in Figure 3 and consists of the following sequential steps:

**Step 1. Initial Contextualization:** A handheld image display begins at the middle slice of the brain or the liver scan. The AI agent introduces general guidance for segmenting the clinical target on the left panel (Figure 3A) using a RAG pipeline. To ground users’ visual understanding, the system retrieves and displays a positive example (with pathology) and a negative example (without pathology) on the right panel, allowing for contrastive comparison. Both examples bear anatomical similarity to the image slice under study. Then, general pathology explanations are delivered via Google’s Text-To-Speech API in

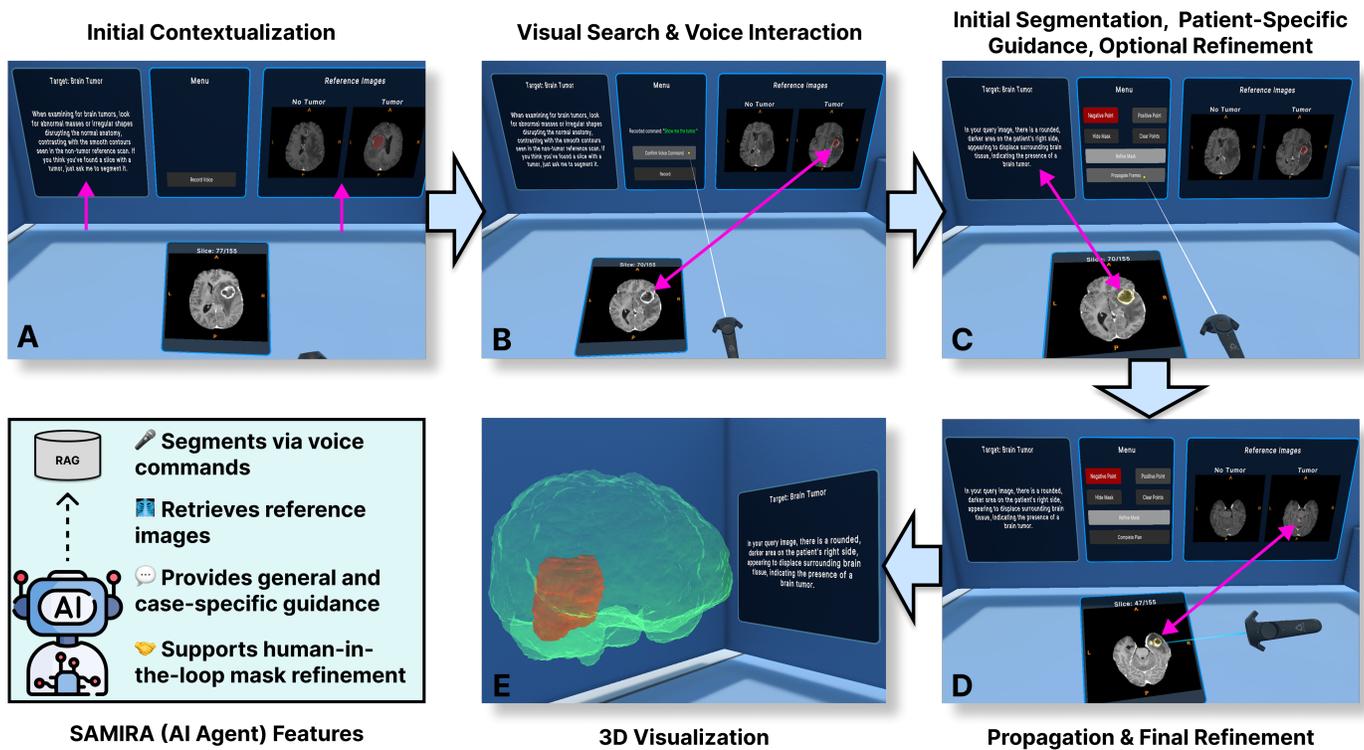


Fig. 3. Demonstration of workflow for the proposed AI-assisted interactive medical image segmentation in VR. Users begin by reviewing AI-generated textual guidance and visually similar reference slices (A), then navigate the volume to find the tumor. Once found, they issue a voice command to segment the tumor (B). Next, the agent predicts a mask and a patient-specific description of the tumor (C). If necessary, users can edit this mask, then propagate it across frames. Finally, users review all predicted masks and place point prompts to refine the masks (D). Upon completion, the final segmented structure is rendered in true 3D scale over the patient’s anatomy (E).

a female voice, and are displayed as text on the left panel. At this stage, the middle panel has one button that says “Record Voice”.

**Step 2. Visual Search:** Users scroll through the image slices in search of the target region and the agent continuously updates the right panel reference panel with visually similar tumor and non-tumor cases retrieved from other patients. The contextual guidance in Step 1 helps less experienced users with deciding whether the current slice contains the pathology of interest (see Figure 3B).

**Step 3. Voice Interaction:** Users initiate segmentation by clicking the “Record Voice” button and issuing natural language requests to the AI agent. Spoken input is transcribed in real time using the Microsoft Azure Speech SDK, with recognized commands displayed on the middle panel (Figure 3B). After voice recognition, users can press the “Record Voice” button again to re-record, or press “Confirm Voice Command” to pass their command to the inference server.

**Step 4. Initial Segmentation and Patient-Specific Guidance:** The AI agent returns an initial segmentation mask at the selected 2D slice and uses its RAG pipeline to provide additional spoken, case-specific diagnostic context (Figure 3C). The middle panel updates with the same refinement options as the *Interaction Paradigm Evaluation* interface and an additional button “Propagate Frames”.

**Step 5. Refinement and Propagation:** Image slice scrolling is disabled at this stage, forcing the users to focus on the

selected 2D slice of interest to achieve the best segmentation accuracy via point-prompt-based refinement if necessary. This is because it will be used to seed the automatic SAM2-based label propagation to obtain the full 3D segmentation. Once satisfied with the current slice, the users initiate mask propagation across slices by pressing the “Propagate Frames” button. As the inference server returns produced segmentation masks in sequence, the image display automatically updates the slices in the order that the masks are produced in real-time, enabling users to visually assess the segmentation results and quickly judge whether additional refinements will be needed.

**Step 6. Final Refinement of Predicted Masks:** Users review the propagated masks across slices, applying additional point-prompt-based refinements wherever segmentation errors remain (Figure 3D). During this stage, the “Propagate Frames” button is replaced by a “Complete Plan” button. When pressed, a confirmation screen appears to prevent accidental submission, ensuring that users have fully completed their corrections before finalizing the segmentation plan. Upon completion, the final segmentation masks are saved.

**Step 7. 3D Visualization:** Upon task completion, a request is sent to the local Python-based websocket server with the path to the saved segmentation result. The server uses the *Marching Cubes* algorithm to extract a polygonal mesh from the 3D segmentation and generates a corresponding .obj file representing the tumor pathology. To ensure anatomically accurate scaling, the voxel resolutions are extracted from the

medical image’s metadata and applied during mesh generation. The resulting mesh is then loaded back into Unity at runtime and rendered at true-to-life scale (Figure 3E). For brain tumor visualization, a threshold of 500 is applied to extract the brain surface from the scan, while liver tumor visualization uses a threshold of 150 to capture relevant anatomical structures such as the spine and ribcage. Additionally, liver tissue from the same patient was rendered using data from the LiTS dataset. In the future, it can easily be generated with liver segmentation models, or even interactively with our workflow.

### C. SAMIRA - Segmentation Algorithm and Retrieval-Augmented Generation

SAMIRA leverages AI foundation models and Retrieval-Augmented Generation to assist the designated interactive medical image segmentation task.

#### 1) Interactive segmentation with foundation models:

We proposed a novel deep learning-based, speech-initiated interactive segmentation method for 3D medical images as illustrated in Figure 4. This method forms a key function of SAMIRA and relies on two complementary foundation models. The first model, *BiomedParse*, generates the initial segmentation mask in response to user-issued voice commands. Spoken prompts are first transcribed using the Azure Speech AI service, producing a text prompt that is passed to *BiomedParse* along with the image slice. Upon receiving the textual prompt and corresponding image slice, *BiomedParse* produces a binary mask of the target structure described in the user’s voice command. The second model, SAM2, is used for interactive refinement and multi-slice propagation. Users can iteratively correct the mask using positive and/or negative point prompts, which are passed to SAM2 for real-time refinement. Once satisfied, the refined mask is propagated bi-directionally through the volume using a modified version of SAM2’s memory mechanism with a novel propagation termination criterion: if the inter-slice Intersection-over-Union (IoU) between the current mask ( $Mask_t$ ) and previous mask ( $Mask_{t-1}$ ) fell below 0.3, propagation is halted in that direction, under the assumption that mask changes should be relatively smooth between neighboring image slices. Inter-slice IoU measures how much two sequentially predicted binary masks overlap, defined as the ratio between the area of their intersection and the area of their union:

$$IoU = \frac{|Mask_t \cap Mask_{t-1}|}{|Mask_t \cup Mask_{t-1}|}$$

While point-prompt-based revision with SAM2 helps ensure the accuracy of the segmentation, it is desirable to keep the needs at minimum for the efficiency of the workflow. Thus, to gauge the baseline performance of our proposed segmentation method in the absence of prompt-based revision, we conducted a standalone ablation study using the Pretreat-MetsToBrain-Masks [24] and LiTS [25] datasets, by comparing the accuracy of a “*BiomedParse* seeding + SAM2 propagation with IoU-based early stopping” pipeline versus a “*BiomedParse* seeding + original SAM2 propagation” one. To accomplish this, we randomly selected 40 volumes (20 brain tumor MRI scans and 20 liver tumor CT scans) from the datasets, excluding cases

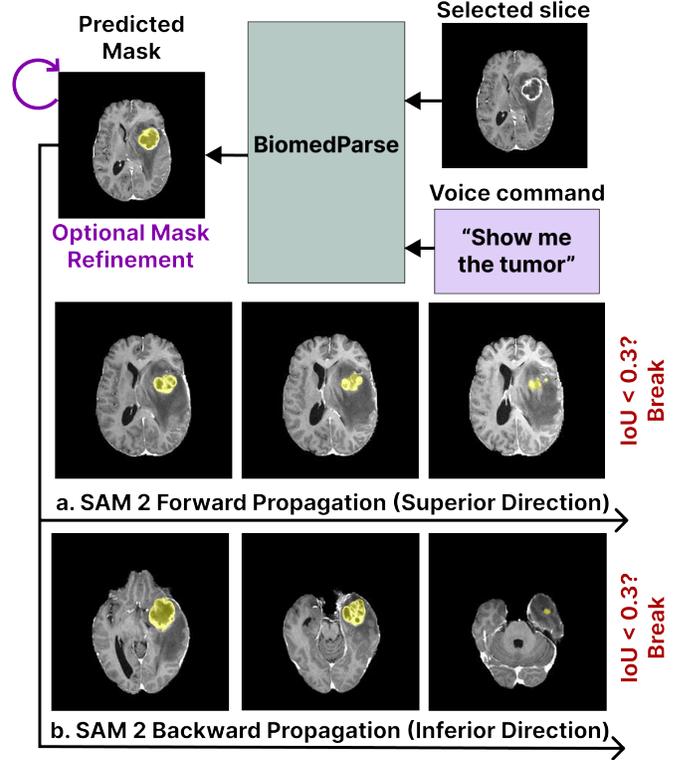


Fig. 4. SAMIRA’s segmentation module for mask prediction, refinement, and propagation across frames. After a voice command initiates initial tumor segmentation via *BiomedParse*, the user may optionally refine the mask through point prompts. The mask is then propagated slice-wise using SAM2, first superiorly (a) and then inferiorly (b), with propagation automatically terminating when inter-slice Intersection-over-Union (IoU) falls below 0.3 to prevent segmentation drift.

used in the full workflow study. For each volume, a slice with visible tumor was manually selected and submitted to the inference server along with a natural language prompt —“show me the brain tumor” for brain cases and “show me the liver tumor” for liver cases. Here, *BiomedParse* generates an initial tumor mask for the selected slice, which was then propagated bi-directionally (in the superior and inferior directions) using SAM2. Using the paired sample Wilcoxon signed-rank test, results indicate that the break condition significantly improved accuracy for liver tumor CT scans ( $p = 0.0024 < 0.05$ ), while having a smaller but still significant effect for brain tumor MRIs ( $p = 0.0039 < 0.05$ ), as summarized in Table I. The stronger effect observed in CT scans may be attributed to the higher levels of noise and lower soft-tissue contrast, which can cause SAM2, originally trained on natural images, to mistake noise for anatomical structures. The break condition prevents propagation of spurious segmentations across slices, which is particularly helpful in noisier CTs. Compared to recent SAM2-based methods for 3D medical segmentation [11], our approach introduces both a language-based initialization and a propagation stopping rule, enhancing accuracy and reducing user burden. These findings support the inclusion of the IoU threshold as an essential mechanism for robust slice-wise mask propagation in clinical datasets.

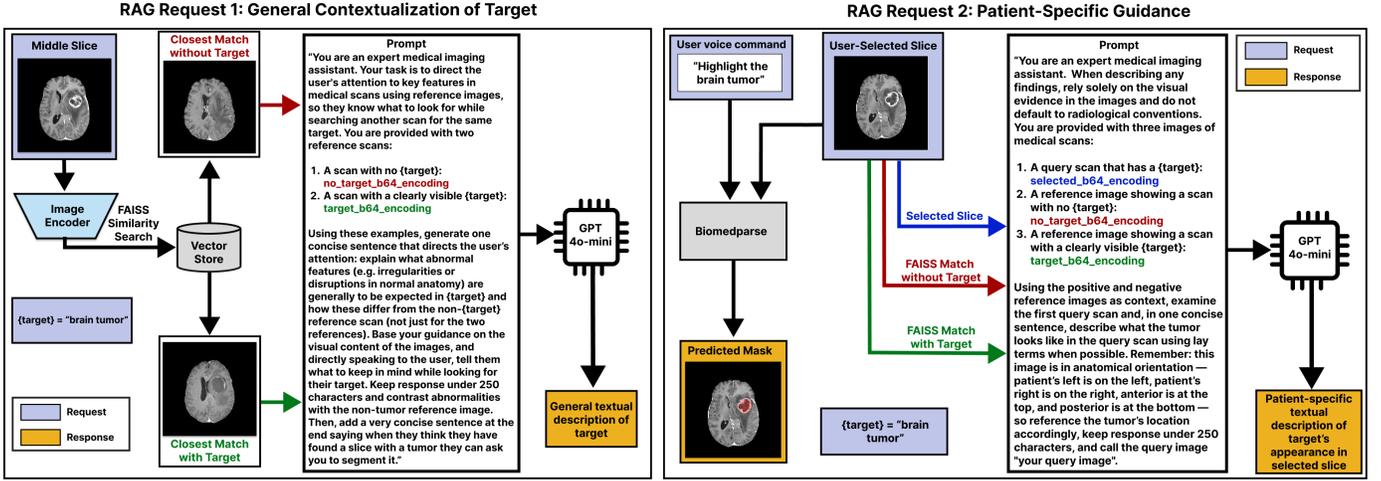


Fig. 5. Retrieval-Augmented Generation (RAG) pipelines for multimodal guidance. (Left) To support initial understanding, the system retrieves two anatomically similar reference slices—one with and one without the target pathology—and uses them to generate a general description of the abnormality. (Right) After the user selects a slice and issues a voice command, the system compares visual features of the patient’s scan to healthy and pathological reference images. Guided by shared and differing features, the agent describes what the abnormality likely looks like in the selected slice.

TABLE I  
MEAN 3D DICE SCORES FOR AUTOMATIC SAM2 PROPAGATION WITH AND WITHOUT THE IOU BREAK CONDITION.

Modality	Target	IoU Break	3D Dice (% , Mean $\pm$ Std)
MRI	Brain Tumor	True	87.41 $\pm$ 10.77 <sup>†</sup>
		False	87.28 $\pm$ 10.73
CT	Liver Tumor	True	73.31 $\pm$ 13.04 <sup>†</sup>
		False	68.94 $\pm$ 16.30

<sup>†</sup>Statistically significant improvement compared to no IoU break condition

2) **RAG-based guidance system:** In addition to interactive segmentation, to provide multi-modal, context-aware guidance during segmentation, SAMIRA employs a RAG framework that integrates image similarity search with generative language modeling. At its core is FAISS [26], a library developed by Meta for fast approximate nearest-neighbor search on high-dimensional vectors. FAISS enables real-time retrieval of anatomically similar tumor and non-tumor slices from large medical image datasets, based on high-dimensional vectors from the high-level feature maps output by the Res5 layer of BiomedParse’s image encoder. These retrieved examples are used in two ways: first, to provide general contextual grounding using representative healthy and pathological slices (RAG Request 1) and second, to generate query-specific guidance based on the user’s spoken prompt and the current image slice (RAG Request 2).

**Knowledge Database Construction:** We constructed two FAISS databases, one for brain MRIs and one for liver CTs, which contain a total of 30,845 brain slices (6,766 tumor, 24,079 without tumor) and 57,193 liver slices (6,982 tumor, 50,210 without tumor), drawn from 199 and 127 patients, respectively. Embedding vectors were computed using BiomedParse, normalized, and stored for efficient similarity search. These databases can easily be expanded in the future with a larger variety of targets and cases.

**RAG Request 1. Initial Contextualization:** When a new scan of interest is loaded in the system, SAMIRA uses the encoding vector of the middle slice to retrieve two visually similar reference images from the knowledge database: one with and one without the target pathology (i.e., tumor). These contrastive examples help ground the task as the non-tumor slice illustrates normal anatomy, while the tumor case highlights typical features of the target pathology. Then, both are passed to GPT-4o-mini, chosen for its fast inference time, which generates a general explanation of what to look for during visual exploration of the target. Figure 5 illustrates this process, showing how supplying reference images with a carefully engineered prompt can support contextual grounding to guide the user.

**RAG Request 2. Query-Specific Interpretation:** While general contextualization offers users an overview of what the pathology typically looks like, it does not account for anatomical variation in individual cases being investigated. Therefore, query-specific interpretation is used to provide more personalized, patient-specific guidance. When the user issues a voice command to segment a structure (e.g., “highlight the brain tumor”), the selected image slice is used to query the knowledge database, retrieving two visually similar reference slices: one with the target pathology (positive example) and one without (negative example). These references, *along with the user’s query slice*, are passed to GPT-4o-mini in another structured multi-modal prompt, illustrated in Figure 5. By explicitly contrasting the *query* with both healthy and pathological references, the agent can generate targeted explanations grounded in the anatomy of the target image under study.

#### D. User Study and Evaluation Metrics

We conducted two separate user studies for the developed system. While the first study aimed to identify the optimal user interaction paradigm for SAM2-based segmentation refinement, the second study assessed the full workflow of

the proposed system. All user studies were conducted under institutional ethics approval and after all participants provided informed consent.

1) *User-Interaction Paradigm Evaluation*: We recruited 15 participants (age =  $26.5 \pm 2.68$  years; 6 female & 9 male) to evaluate the three interaction paradigms (*Controller*, *Head Pointing*, and *Eye Tracking*). Participants rated their familiarity with VR, human anatomy, medical imaging modalities, and medical image segmentation on a 1–5 scale (1 = unfamiliar, 5 = familiar). On average, they reported familiarity with VR ( $4.2 \pm 0.94$ ), and between neutral and somewhat familiarity with anatomy ( $3.47 \pm 1.24$ ), imaging modalities ( $3.53 \pm 1.30$ ), and medical image segmentation ( $3.73 \pm 1.27$ ). Prior to the study, participants completed a guided hands-on tutorial in the VR environment, which included practice with segmentation mask refinement on a lung CT scan from the LCTSC dataset [27] and menu interaction using each of the three interaction paradigms. The tutorial session also served to calibrate the eye-tracking system. Following this, users began the interaction test, correcting the erroneous BRATS brain tumor MRI slices with each interaction paradigm as described in Section IV.A. For each interaction paradigm trial, task completion time, segmentation accuracy (with 3D Dice score), perceived task load, point prompts placed, and point prompts erased were recorded.

To assess perceived task load, we used the *NASA Task Load Index (NASA-TLX)* [28], a validated six-item instrument measuring mental demand, physical demand, temporal demand, effort, frustration, and perceived performance. Each item was rated in the range of 1–21, which was scaled to 0–100 for analysis, and finally all items were averaged to produce a total task load score. To compare overall performance across interaction paradigms, we computed a composite interaction score for each of the 45 user-paradigm trials (15 participants  $\times$  3 interaction paradigms). This score quantifies trade-offs between segmentation accuracy, task load, and completion time. Z-scores for accuracy, NASA-TLX, and completion time were calculated across all trials to ensure standardized comparison. For each trial  $i$ , the composite score was calculated as:

$$\text{Composite}_i = z_{\text{accuracy}} - z_{\text{NASA}} - z_{\text{time}}$$

This score treats accuracy as beneficial and both task load and time as costs, giving equal weight to each. Mean and standard deviation of composite scores were then computed per interaction paradigm to summarize performance. Finally, for additional self-reported confirmation, we asked users to rank their preferred interaction paradigm.

2) *Full Workflow Study of SAMIRA*: To evaluate the complete AI-assisted segmentation system, including conversational interaction, mask refinement, and 3D visualization, we recruited 19 participants (age =  $26.8 \pm 3.63$  years; 8 female & 11 male). On average, they again reported familiarity with VR ( $4.21 \pm 1.18$ ) and between neutral to somewhat familiarity with anatomy ( $3.42 \pm 1.22$ ), imaging modalities ( $3.79 \pm 1.18$ ), and medical image segmentation ( $3.74 \pm 1.37$ ). After a brief tutorial, where participants practiced the full workflow by segmenting a lung CT scan from the LCTSC

dataset [27], participants completed the two segmentation tasks: brain tumor in MRI and liver tumor in CT. The order of these was randomized across participants to minimize order effects. Outcome measures included segmentation accuracy (3D Dice scores before and after segmentation refinement), task completion time, and user experience metrics. Task load was assessed using NASA-TLX, similarly to the interaction paradigm evaluation. To assess the overall usability, we also administered the *System Usability Scale (SUS)* [29], a widely used 10-item questionnaire that yields a total score from 0 to 100. Scores above 68 are considered to indicate good usability [30]. Additionally, we developed a custom 9-item questionnaire to evaluate users’ perceptions of AI agent guidance, reference images, 3D visualization, and user confidence in the workflow. The full list of questions are provided in Figure 7. Items used a 5-point Likert scale from 1 (strongly disagree) to 5 (strongly agree). Finally, we included an open-ended section, where users could freely mention what they liked and/or disliked about the system, and further elaborate their semi-quantitative evaluations.

3) *Statistical Analysis*: For the interaction paradigm study, differences in NASA-TLX scores, completion times, and Dice scores across three user-interaction paradigms were evaluated using Kruskal–Wallis tests. For the full workflow study, we compared the 3D Dice scores before and after user refinement using Wilcoxon rank-sum tests. SUS scores were tested against the usability benchmark of 68 using a one-sample  $t$ -test. Custom user questionnaire responses were tested against the neutral value of 3 using Wilcoxon signed rank tests. For all three statistical tests, a statistical significance was confirmed with  $p < 0.05$ .

## V. RESULTS

### A. Interaction Paradigm Evaluation

TABLE II  
COMPARISON OF INTERACTION PARADIGMS FOR SEGMENTATION REFINEMENT. VALUES ARE SHOWN AS MEAN  $\pm$  STANDARD DEVIATION. THE BEST SCORE IS IN BOLD FONTS. NASA-TLX IS OUT OF 100.

Interaction Mode	3D Dice (%)	Time (s)	NASA-TLX	Composite Score
Controller	<b>99.25 <math>\pm</math> 00.25</b>	<b>220.3 <math>\pm</math> 79.3</b>	18.8 $\pm$ 14.5	<b>0.51 <math>\pm</math> 1.91</b>
Head Pointing	99.21 $\pm$ 0.30	248.8 $\pm$ 78.5	<b>16.8 <math>\pm</math> 13.9</b>	0.20 $\pm$ 1.56
Eye Tracking	99.13 $\pm$ 0.46	251.1 $\pm$ 78.7	26.6 $\pm$ 15.4	-0.71 $\pm$ 1.77

1) *Accuracy, Completion Time, and NASA-TLX*: The metrics for evaluating the three interaction paradigms are shown in Table II, with non-significant differences in overall scores between groups ( $p > 0.05$ ). For **3D Dice (segmentation accuracy)**, all paradigms yielded highly accurate refined masks, significantly above the starting 3D Dice score of 0.91 ( $p < 0.05$ ). The *Controller* paradigm had the highest accuracy score with the least variability ( $99.25 \pm 0.25\%$ ), followed by *Head Pointing* ( $99.21 \pm 0.30\%$ ), then *Eye Tracking* ( $99.13 \pm 0.46\%$ ). For **completion time**, the ranking remained the same: *Controller* was the fastest ( $220.3 \pm 79.3s$ ) followed by *Head Pointing* ( $248.8 \pm 78.5s$ ) and *Eye Tracking* ( $251.1 \pm 78.7s$ ). The overall **NASA-TLX** scores (out of 100) slightly favored *Head Pointing*, which had the lowest overall

task load ( $16.8 \pm 13.9$ ), followed by *Controller* ( $18.8 \pm 14.5$ ) and *Eye Tracking* ( $26.6 \pm 15.4$ ). While overall task load did not demonstrate significant differences, one sub-item showed a significant effect: mental demand was significantly lower for *Head Pointing* ( $16.7 \pm 22.5$ ) compared to *Controller* ( $30.0 \pm 23.9$ ,  $p = 0.0123$ ) and *Eye Tracking* ( $38.7 \pm 25.5$ ,  $p = 0.0329$ ), as illustrated in Figure 6. No other NASA-TLX sub-items’ differences reached statistical significance.

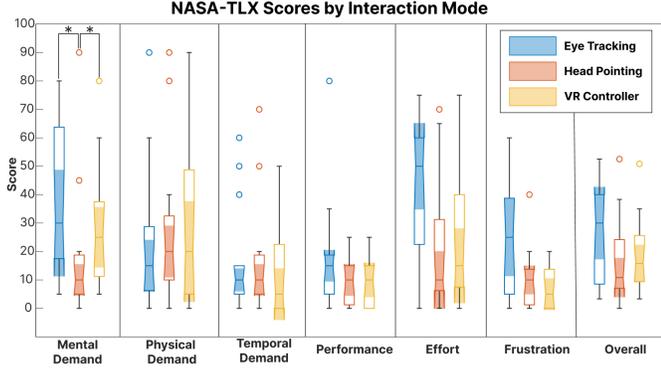


Fig. 6. Boxplots of NASA Task Load Index (TLX) for different interaction paradigms, including Controller, Head Pointing, and Eye tracking. Significant pair-wise differences are marked with “\*\*\*”.

2) **Composite Scores:** The resulting mean composite scores that promote a balance between performance, speed, and user effort were  $0.51 \pm 1.91$  for *Controller*,  $0.20 \pm 1.56$  for *Head Pointing*, and  $-0.71 \pm 1.77$  for *Eye Tracking*. Among the paradigms, *Controller* achieved the highest mean composite score, driven by a favorable combination of high segmentation accuracy (99.25%), low task load (NASA-TLX = 18.8), and the fastest completion time (220.3 s). *Head Pointing* followed closely, while *Eye Tracking* lagged behind due to relatively higher task load and slower performance.

3) **Point Prompt Efficiency:** To complement the accuracy, task completion time, and NASA-TLX, we also compared the three interaction paradigms in terms of the ratio between the number of confirmed point prompts for segmentation refinement and the cleared points (i.e., points-per-clear efficiency, the higher the better). This ratio reveals the robustness of these paradigms with the joint consideration of their inherent precision, efficiency, and task load. Overall, *Controller* yielded the highest points-per-clear efficiency ( $24.33 \pm 24.91$ ), followed by *Head Pointing* ( $16.08 \pm 6.17$ ) and *Eye tracking* ( $14.67 \pm 8.72$ ), suggesting that the *Controller* condition was more robust for point prompt placement (less corrections required) and eye tracking was more prone to errors.

4) **User preferences:** Finally, when asking the participants to rank their preferred interaction paradigms, 7 of 15 participants selected *Controller*, 6 selected *Head Pointing*, and only 2 selected *Eye tracking* as their top choices. These results, in addition to the close composite scores, suggest that both *Controller* and *Head Pointing* are viable and well-received paradigms, while *Eye tracking* is comparatively less favored. The close ranking between *Controller* and *Head Pointing* indicates that either approach could be suitable for deployment in the full workflow. However, to reduce variability in

downstream evaluation, we selected *Controller* as the primary interaction paradigm for the subsequent full workflow study due to its marginally superior results.

## B. Full Segmentation Workflow with SAMIRA

1) **Segmentation accuracy:** In terms of segmentation accuracy (3D Dice), the workflow yielded high scores. For the **brain MRI data**, refined masks ( $94.92 \pm 0.52\%$ ) showed significantly higher accuracy than the user’s propagated unrefined masks ( $90.53 \pm 10.51\%$ ,  $p < 0.0001$ ). For the **liver CT data**, Dice scores were comparable between unrefined ( $95.47 \pm 0.39$ ) and refined ( $95.46 \pm 0.40$ ) masks, with no significant difference, ultimately indicating high starting accuracies are preserved (Table III). This high starting accuracy is likely attributed to the liver tumor case being larger and more spherical shaped than the brain tumor case that appears as a mass centrally, but splits into multiple lobes in the superior and inferior regions. The lower difficulty of the liver tumor case is further reflected in the shorter completion time (liver:  $279.6 \pm 109.4$ s vs. brain:  $361.7 \pm 144.3$  s) and fewer point prompts placed (liver:  $8.8 \pm 6.5$  vs. brain:  $27.7 \pm 19.9$ ). The brain tumor case likely demanded more user input due to its aforementioned irregular shape.

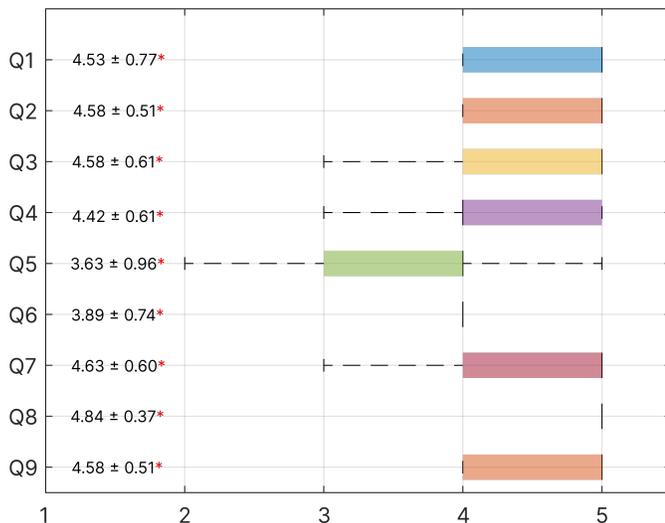
2) **Semi-quantitative questionnaire results:** Participants rated the overall workflow with SAMIRA as highly usable. The **System Usability Scale (SUS)** score was  $90.00 \pm 8.98$ , which is significantly higher than the benchmark of 68 ( $p < 0.001$ ) and corresponds to an ‘A’ usability score [30]. Meanwhile, the **NASA-TLX** scores (out of 100) indicate low to moderate task loads across sub-items. Descriptive statistics are as follows: Mental Demand ( $31.84 \pm 28.10$ ), Physical Demand ( $12.11 \pm 14.27$ ), Temporal Demand ( $13.68 \pm 18.25$ ), Performance ( $20.00 \pm 24.72$ ), Effort ( $25.26 \pm 15.50$ ), Frustration ( $7.11 \pm 8.22$ ), and Overall task load ( $18.33 \pm 11.93$ ).

TABLE III  
3D DICE SCORES BEFORE AND AFTER REFINEMENT. ASTERISKS DENOTE STATISTICALLY SIGNIFICANT CHANGES.

Tumor	Unrefined	Refined	p-value
Brain	$90.53 \pm 10.51$	$94.92 \pm 0.52$	$1.19 \times 10^{-5*}$
Liver	$95.47 \pm 0.39$	$95.46 \pm 0.40$	0.3968

Finally, all responses on the **custom questionnaire** were significantly above the neutral midpoint of 3 ( $p < 0.05$ ), indicating favorable perceptions of the system’s AI agent integration, its ability to support pathology learning, the helpfulness of reference images, and the usefulness of 3D visualization (see Figure 7).

3) **Qualitative feedback results:** Participants’ written feedback, where they were asked to describe what they liked or disliked about the system, further supports the system’s perceived usability and educational value. 16 out of 19 participants described the system as “easy to use,” “intuitive,” or “clear,” highlighting its low learning curve and smooth integration with the VR environment. One participant noted, “After you learn it, it is easy to use, fast, and interactive,” while another stated, “The system is very user friendly and well integrated with the HTC Vive controls.” These perceptions align with the high



Q1. I found the system provides an immersive experience.  
 Q2. I found that the AI system is nicely integrated in the workflow.  
 Q3. I found the AI system instrumental for me to complete the segmentation tasks.  
 Q4. Overall, the system offers me a good understanding of the pathology to segment.  
 Q5. I found the reference images helpful in making my judgement for segmentation.  
 Q6. I found the AI-generated description of pathology helpful for me.  
 Q7. I found the 3D visualization of the segmentation result helpful.  
 Q8. Overall, I think the workflow is well designed.  
 Q9. Overall, I think the system improves my confidence for performing medical image segmentation tasks.

Fig. 7. Boxplots of custom user experience questionnaire results (mean  $\pm$  std on left, values significantly above 3 with red asterisk) for the full workflow with SAMIRA.

overall scores on SUS Question 3 (“I thought the system was easy to use,”  $4.58 \pm 0.51$ ) and Question 5 (“I found the various functions in this system were well-integrated,”  $4.68 \pm 0.48$ ). AI voice interaction was also frequently praised. 10 participants explicitly mentioned that the voice command feature improved usability, with one stating, “The voice commands are well integrated and the propagation is very helpful for identifying the whole tumor.” Another remarked, “The AI was well integrated and didn’t feel intrusive—more like a helpful assistant.” The system was also perceived as supportive of learning and decision-making. 7 participants reported increased confidence during segmentation tasks, attributing this to guidance from the AI agent and the real-time visualization tools. For example, one wrote, “It was so useful and made me confident to do the task”, echoing high SUS Question 9 results (“I felt very confident using this system,”  $4.74 \pm 0.45$ ).

While overall impressions were positive, some participants suggested providing additional medical context, and a few expressed awkwardness with voice inputs. Nonetheless, the qualitative feedback reinforces the questionnaire findings and confirms that SAMIRA’s design successfully balances guidance, autonomy, and interpretability in a VR environment.

## VI. DISCUSSION

The findings of our experiment provided partial support for our hypothesis that *Head Pointing* would provide the optimal trade-off between segmentation accuracy, efficiency, and task load. While *Head Pointing* did show the lowest mental demand, *Controller*-based input achieved slightly better

overall performance, as reflected in its higher mean accuracy and efficiency metrics and its highest average composite score. Furthermore, all paradigms yielded excellent Dice scores following segmentation correction, but *Controller* and *Head Pointing* outperformed *Eye Tracking* in terms of accuracy, task load, and completion time, echoing the findings of Xu et al.’s evaluation of text-selection techniques [20], where head-pointing and controller performance were close. Despite composite scores favoring controller-based pointing overall, *Head Pointing* emerged as a lower mental effort alternative—especially for applications where users may need to work with just one hand or seek a cognitively lighter interaction. *Eye Tracking*, while promising in theory, remains less favored for segmentation refinement tasks, where precision and visual stability are critical, despite the damping function (Equation 1) we employed to improve precision and user-experience. In future deployments, the system can allow users to select their preferred interaction paradigm, offering flexibility. This flexibility would be warranted, since all three paradigms were chosen as favorites across the different users.

Overall, our findings suggest that AI-assisted segmentation in VR is not only technically viable, but also educationally and ergonomically impactful. Across studies, users were able to generate high-quality segmentations without domain expertise and the system demonstrated that it promotes understanding, confidence, and informed interaction. In the full workflow, participants rated the system highly on both usability and interpretability of the task. As seen in Figure 7, users strongly agreed that the AI agent helped them complete tasks (Q3:  $4.58 \pm 0.61$ ), supported understanding of the pathology (Q4:  $4.42 \pm 0.61$ ), and improved their confidence in performing segmentation tasks (Q9:  $4.58 \pm 0.51$ ). The RAG mechanism played a key role here. By comparing the queried radiological slice to both healthy and pathological references, users received contextualized, anatomy-specific guidance that was grounded in real cases. This was especially reflected in high agreement with Q8, where workflow design was rated highest ( $4.84 \pm 0.37$ ). Interestingly, the AI-generated textual explanations (Q6) and reference images (Q5) received slightly lower scores ( $3.89 \pm 0.74$  and  $3.63 \pm 0.96$ , respectively), possibly reflecting uncertainty and confusion in some participants, who in general did not have strong familiarity with human anatomy ( $3.42 \pm 1.22$ , 1 = unfamiliar, 5 = familiar). To improve comfort and presence during voice interaction, future versions of the system could feature a visual avatar for the assistant, helping to reduce the slight awkwardness some users felt when speaking to a disembodied voice.

Furthermore, it was evident in the segmentation results that users demonstrated a clear understanding of when to intervene and when to trust the AI system. For example, with the simpler liver tumor CT case, users made minimal edits, and the refined masks were similar to unrefined ones. Yet importantly, performance did not degrade after user interaction, indicating that users did not over-correct or introduce noise. This suggests a healthy level of trust and restraint, and a true understanding of segmentation quality based on visual features and reference images. In contrast, the more challenging brain tumor MRI case, which included branching outer boundaries, showed a

significant accuracy improvement after refinement. The 3D Dice score rose from  $90.53 \pm 10.51\%$  to  $94.92 \pm 0.52\%$  post-interaction ( $p < 0.001$ ), demonstrating that users could identify areas needing correction and effectively apply point-based refinements. This reinforces that users not only understand segmentation correctness, but can also meaningfully enhance AI outputs in cases where human expertise is still necessary.

Together, these results position our system not only as a segmentation tool, but as a supportive assistant capable of accelerating workflows, teaching radiological features of pathology, and fostering trust with clinical AI. The positive responses to confidence and task understanding, combined with high segmentation accuracy, suggest that such a tool has potential in both clinical workflows and medical education. Future work may explore expanding the agent to work with more radiological concepts, or adaptive RAG responses based on user skill level (i.e. beginners versus experienced radiologists).

## VII. CONCLUSION

We introduced a novel VR system for interactive medical image segmentation that integrates foundation models with attention-switching interaction and a supportive conversational AI agent, SAMIRA. At the core of our method is a novel segmentation algorithm that combines BiomedParse’s language-driven detection with a medical-image-adapted SAM2 model. To adapt SAM2 for clinical imaging, we introduced a novel IoU-based stopping criterion in its memory mechanism to prevent drift across noisy or low-contrast slices. Our findings show that this criterion can significantly improve segmentation quality, and that by using it with SAMIRA, users can efficiently achieve high segmentation accuracy with minimal effort across all interaction paradigms. Of these interaction paradigms, *Controller* pointing offers the best overall balance of accuracy, speed, and task load, closely followed by head-pointing. More importantly, users demonstrated a clear ability to interpret and refine AI outputs based on its generated guidance, engaging critically with reference images, contextual explanations, and 3D visualization.

Importantly, this system is generalizable to any 3D medical image aligned with BiomedParse’s training scope and can be expanded by enriching its RAG knowledge database. Its modular design, high usability, and adaptability to various interaction styles position it as a powerful tool for clinical workflows, but also offers insights for future HCI research in intelligent, immersive medical systems with clinical AI agents.

## ACKNOWLEDGMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (596537, RGPIN/05100-2022), the Fonds de recherche du Québec – Nature et technologies (B1X-348625), and the Fonds de recherche du Québec – Santé Junior 1 Research Scholar program. We thank Taha Koleilat for the insightful discussion on medical foundation models.

## REFERENCES

- [1] M. A. Bruno, E. A. Walker, and H. H. Abujudeh, “Understanding and confronting our mistakes: the epidemiology of error in radiology and strategies for error reduction,” *Radiographics*, vol. 35, no. 6, pp. 1668–1676, 2015.
- [2] M. Konenkov, A. Lykov, D. Trinitatova, and D. Tsetserukou, “Vr-gpt: Visual language model for intelligent virtual reality applications,” *arXiv preprint arXiv:2405.11537*, 2024.
- [3] M. Behravan, K. Matković, and D. Gračanin, “Generative ai for context-aware 3d object creation using vision-language models in augmented reality,” in *2025 IEEE International Conference on Artificial Intelligence and eXtended and Virtual Reality (AIxVR)*. IEEE, 2025, pp. 73–81.
- [4] S. Budd, E. C. Robinson, and B. Kainz, “A survey on active learning and human-in-the-loop deep learning for medical image analysis,” *Medical image analysis*, vol. 71, p. 102062, 2021.
- [5] C. Vásquez-Venegas, C. G. Sotomayor, B. Ramos, V. Castañeda, G. Pereira, G. Cabrera-Vives, and S. Härtel, “Human-in-the-loop—a deep learning strategy in combination with a patient-specific gaussian mixture model leads to the fast characterization of volumetric ground-glass opacity and consolidation in the computed tomography scans of covid-19 patients,” *Journal of clinical medicine*, vol. 13, no. 17, p. 5231, 2024.
- [6] X. Jin, Y. Hao, J. Hilliard, Z. Zhang, M. A. Thomas, H. Li, A. K. Jha, and G. D. Hugo, “A quality assurance framework for routine monitoring of deep learning cardiac substructure computed tomography segmentation models in radiotherapy,” *Medical physics*, vol. 51, no. 4, pp. 2741–2758, 2024.
- [7] S. Y. Liaw, J. Z. Tan, K. D. Bin Rusli, R. Ratan, W. Zhou, S. Lim, T. C. Lau, B. Seah, and W. L. Chua, “Artificial intelligence versus human-controlled doctor in virtual reality simulation for sepsis team training: randomized controlled study,” *Journal of medical Internet research*, vol. 25, p. e47748, 2023.
- [8] V. Chheang, S. Sharmin, R. Márquez-Hernández *et al.*, “Towards anatomy education with generative ai-based virtual assistants in immersive virtual reality environments,” in *Proc. IEEE Int. Conf. on Artificial Intelligence and Virtual Reality (AIVR)*, 2024, pp. 21–30.
- [9] T. Zhao, Y. Gu, J. Yang *et al.*, “Biomedparse: a biomedical foundation model for image parsing of everything everywhere all at once,” *Nature Methods*, vol. 22, pp. 166–176, 2025.
- [10] A. Kirillov *et al.*, “Sam 2: Segment anything, any time, all at once,” *arXiv preprint arXiv:2404.07143*, 2024.
- [11] C. Shen, W. Li, Y. Shi, and X. Wang, “Interactive 3d medical image segmentation with sam 2,” 2025. [Online]. Available: <https://arxiv.org/abs/2408.02635>
- [12] J. Zhu, A. Hamdi, Y. Qi, Y. Jin, and J. Wu, “Medical sam 2: Segment medical images as video via segment anything model 2,” *arXiv preprint arXiv:2408.00874*, 2024.
- [13] Y. Gao, C. Chang, X. Yu, P. Pang, N. Xiong, and C. Huang, “A vr-based volumetric medical image segmentation and visualization system with natural human interaction,” *Virtual Reality*, vol. 26, no. 2, pp. 415–424, 2022.
- [14] S. B. Yonker, O. O. Korshak, T. Hedstrom, A. Wu, S. Atre, and J. P. Schulze, “3d medical image segmentation in virtual reality,” *Electronic Imaging*, vol. 31, pp. 1–6, 2019.
- [15] S. González Izard, R. Sánchez Torres, O. Alonso Plaza, J. A. Juanes Mendez, and F. J. García-Peñalvo, “Nextmed: automatic imaging segmentation, 3d reconstruction, and 3d model visualization platform using augmented and virtual reality,” *Sensors*, vol. 20, no. 10, p. 2962, 2020.
- [16] U. Rashid, J. Kauko, J. Häkkinen, and A. Quigley, “Proximal and distal selection of widgets: designing distributed ui for mobile interaction with large display,” in *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, 2011, pp. 495–498.
- [17] L. Sidenmark, F. Prummer, J. Newn, and H. Gellersen, “Comparing gaze, head and controller selection of dynamically revealed targets in head-mounted displays,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 11, pp. 4740–4750, 2023.
- [18] F. L. Luro and V. Sundstedt, “A comparative study of eye tracking and hand controller for aiming tasks in virtual reality,” in *Proceedings of the 11th ACM Symposium on eye tracking research & applications*, 2019, pp. 1–9.
- [19] A. S. Fernandes, T. S. Murdison, and M. J. Proulx, “Looking in depth: Targeting by eye and controller input for multi-depth target placement,” *International Journal of Human-Computer Interaction*, pp. 1–16, 2024.

- [20] W. Xu, X. Meng, K. Yu, S. Sarcar, and H.-N. Liang, "Evaluation of text selection techniques in virtual reality head-mounted displays," in *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2022, pp. 131–140.
- [21] O. Hellum, M. Kersten-Oertel, and Y. Xiao, "Assessment of user-interaction strategies for neurosurgical data navigation and annotation in virtual reality," *Virtual Reality*, vol. 27, no. 2, pp. 1345–1355, 2023.
- [22] A. K. Mutasim, A. U. Batmaz, and W. Stuerzlinger, "Pinch, click, or dwell: Comparing different selection techniques for eye-gaze-based pointing in virtual reality," in *Acm symposium on eye tracking research and applications*, 2021, pp. 1–7.
- [23] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest *et al.*, "The multimodal brain tumor image segmentation benchmark (brats)," *IEEE transactions on medical imaging*, vol. 34, no. 10, pp. 1993–2024, 2014.
- [24] D. Ramakrishnan, L. Jekel, S. Chadha, A. Janas, H. Moy, N. Maleki, M. Sala, M. Kaur, G. C. Petersen, S. Merkaç *et al.*, "A large open access dataset of brain metastasis 3d segmentations on mri with clinical and imaging information," *Scientific Data*, vol. 11, no. 1, p. 254, 2024.
- [25] P. Bilic, P. Christ, H. B. Li, E. Vorontsov, A. Ben-Cohen, G. Kaissis, A. Szeskin, C. Jacobs, G. E. H. Mamani, G. Chartrand *et al.*, "The liver tumor segmentation benchmark (lits)," *Medical image analysis*, vol. 84, p. 102680, 2023.
- [26] M. Douze, A. Guzhva, C. Deng, J. Johnson, G. Szilvasy, P.-E. Mazaré, M. Lomeli, L. Hosseini, and H. Jégou, "The faiss library," *arXiv preprint arXiv:2401.08281*, 2024.
- [27] J. Yang, G. Sharp, H. Veeraraghavan, W. Van Elmpt, A. Dekker, T. Lustberg, and M. Gooding, "Data from lung ct segmentation challenge 2017 (lctsc)," 2017. [Online]. Available: <https://www.cancerimagingarchive.net/collection/lctsc/>
- [28] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.
- [29] J. Brooke *et al.*, "Sus-a quick and dirty usability scale," *Usability evaluation in industry*, vol. 189, no. 194, pp. 4–7, 1996.
- [30] J. Brooke, "Sus: a retrospective." *Journal of usability studies*, vol. 8, no. 2, 2013.