

An interaction metaphor for enhanced VR-based volume segmentation

E. Monclús¹  and P. Vázquez¹ 

¹ViRVIG Group - Universitat Politècnica de Catalunya

Abstract

The segmentation of medical models is a complex and time-intensive process required for both diagnosis and surgical preparation. Despite the advancements in deep learning, neural networks can only automatically segment a limited number of structures, often requiring further validation by a domain expert. In numerous instances, manual segmentation is still necessary. Virtual Reality (VR) technology can enhance the segmentation process by providing improved perception of segmentation outcomes and enabling interactive supervision by experts. But inspecting how the progress of the segmentation algorithm is evolving, and defining new seeds requires seeing the inner layers of the volume, which can be costly and difficult to achieve with typical metaphors such as clipping planes. In this paper, we introduce a wedge-shaped 3D interaction metaphor designed to facilitate VR-based segmentation through detailed inspection and guidance. User evaluations demonstrated increased satisfaction with usability and faster task completion times using the tool.

CCS Concepts

• **Human-centered computing** → *Interaction techniques; Visualization systems and tools*; • **Applied computing** → *Health informatics*;

1. Introduction

Segmentation of medical images is an essential task for diagnosis, pre-surgery training, and other applications. It is a tedious and time-consuming task. And despite researchers have been working to develop automatic methods, results are not perfect, and a significant amount of manual work is still required, either to complete the segmentation of complex models or to verify the results obtained by automatic algorithms. Advanced interaction techniques may improve both quality and time efficiency [LNBK*16]. And VR has emerged as a valid alternative that can be useful in tasks such as surgical preparation [LSC*21]. In the context of segmentation, VR is powerful since it provides an improved perception of structures. However, seed placement, a frequent task, requires selecting points in the volume using medical images as a reference. Simple clipping planes only provide limited context. Multiple clipping planes in a box is another option, but it makes inspecting the inner parts of the volume difficult. Individual, configurable polygons showing the slices of the volume are also used in desktop software (e.g., [PHK04]), but their manipulation is cumbersome.

We have developed an interaction metaphor centered around a wedge-shaped tool that streamlines and expedites segmentation inspection and guidance. Our tool offers several advantages: *a)* enhanced context: Utilizing a wedge configuration with three clip-

ping planes, our tool allows for toggling medical images on and off, as well as adjusting their transparency, *b)* flexible manipulation: The wedge tool can be easily positioned using a controller, and the planes can be rotated or translated, and *c)* Segmentation guidance: The segmentation process can be guided and controlled using a second controller, providing precision. The final design of the tool has been achieved through an iterative process. Initially, a design was formulated featuring the same wedge-shaped form, accompanied by additional methods for executing equivalent manipulations. This design was then subjected to testing with seven users. The results revealed that the initial design was perceived as user-friendly and more efficient than the utilization of a clipping plane. However, users encountered cognitive challenges in understanding the initial controls. Consequently, we proceeded to redesign the interaction mechanisms and incorporate enhanced visual cues to facilitate direct manipulation of the various elements.

2. Related Work

Medical image segmentation has witnessed significant advancements over the years, transitioning from labor-intensive manual approaches to sophisticated neural networks that automate much of the computational burden [WLC*22]. However, despite the efficacy of automatic methods, they are not entirely immune to errors.




Figure 1: The different elements that can be selected to interact with the wedge shape. The cylinders are used to manipulate planes, while the spheres are used to change the medical image configuration (on and off, as well as window and level).

Certain challenges arise, such as the absence of ground truth for numerous anatomical structures or the presence of substantial variability in the obtained results [TPL*22]. Additionally, the segmentation of moving structures, like the colon, poses further complexities [OMN*23]. Many volume rendering techniques commonly rely on clipping planes to explore the internal structures of volumetric models. However, these techniques often lack the necessary flexibility to enable detailed exploration. Consequently, alternative approaches have been proposed, such as various types of lenses and cutaways [TGK*17, SHT19]. It is worth noting that while notable exceptions exist (e.g., [FP00, MVN13, THY*19]), the majority of these approaches have been primarily designed with desktop devices in mind, rather than 3D interaction in virtual reality (VR) environments. This limitation is significant as the precision and other challenges associated with desktop devices differ from those encountered in VR. Particularly, accurate selection in VR, crucial for seed placement in interactive segmentation, poses a unique challenge as maintaining elevated arms for extended periods can be fatiguing [SKM*23]. Consequently, specific techniques have been developed. Previous approaches include enhanced cursors [BPC19, LYS20], additional helper views [MHG*23], or progressive methods [KBB11, MV23].

3. A Wedge Tool for enhanced interaction in VR

Our tool takes the form of a wedge, consisting of three planes denoted as Π_A , Π_B , and Π_C , as depicted in Figure 1, that serve as seed placement points while providing contextual information by projecting the medical data onto them. The tool is generated upon user request and appears anchored to the pointer ray. Upon creation, the initial task for the user is to position the inspection tool at the desired point. During this stage, the tool remains aligned with the pointer, as shown in Figure 2, and the edge collinear with the ray is highlighted, indicating that it will move in response to user interaction. The wedge can be locked using the application menu, but it can still be repositioned by pointing at the wedge and activating the trigger to enable drag based on the ray's position. Once the wedge




Figure 2: Initial creation of the inspection tool. Until it is locked at a position, the wedge-shaped tool follows the ray.

is locked in place, other inspection tasks become available. This design choice is driven by the contrast between our wedge tool and a cubic tool. While the clipping planes of a cubic tool intersect the exterior parts of the volume, focusing on the external surfaces of the cube for inspection, the wedge tool allows for the visualization of its interior, as illustrated in Figure 2.

3.1. Interaction design

The wedge can be manipulated with the following interactions: *a)* Moving and locking the wedge, *b)* rotating and moving the A and B planes, *c)* changing the distance of the C plane, and *d)* toggling on and off the medical images, and changing its window level, window, and alpha values. Our selection system encompasses diverse components, including planes, their boundaries, and specific vertices (see Figure 1). In addition to the planes themselves, the remaining components are represented as cylinders and spheres. When the user's ray points towards these elements, they are highlighted, thus offering visual feedback of their potential activation. Our tool offers multiple free and relative modes of movement. The latter is designed to achieve more precise adjustments, especially when dealing with smaller modifications, where fine-tuning could become challenging. These movements can be accomplished as:

- Free wedge movement: If the ray is intersecting with the wedge, the trigger will move the tool as the ray moves.
- Relative wedge movement: When the ray is not intersecting the tool, by pressing the trigger, the wedge starts a smooth movement towards the controller.

One key aspect of our design is to avoid users having to divert their gaze away from the model. Thus, we implemented our visual feedback directly on the model, to enable direct manipulation. To achieve manipulation of the elements, we have implemented a combination of selection along with buttons, joystick, and trigger inputs. This interaction scheme allows users to select the desired elements and subsequently utilize the buttons, joystick, and trigger

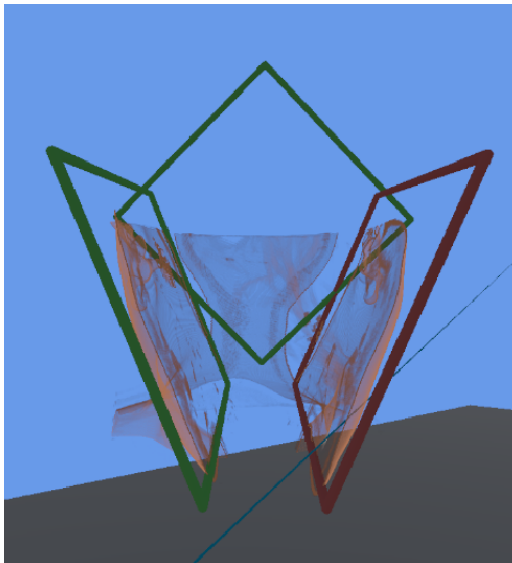


Figure 3: Rotation of the planes.

for effective manipulation. The application menu serves to alternate the wedge's state between locked and unlocked positions. When the wedge is locked, a range of manipulations becomes accessible. The targeted object that necessitates selection, and its corresponding controls and manipulations are:

- Angle cylinder: Trigger + controller movements changes the angle of Π_A or Π_B planes.
- Distance cylinder: Trigger + controller movements changes the distance of the Π_C plane.
- Plane Π_A or Π_B + joystick: up and down moves the plane against or in the direction of the plane normal, respectively
- Wedge + controller movement drags the wedge.
- On-off sphere + trigger toggles medical data on and off.
- Window level sphere + trigger + joystick: left and right will increase and decrease the window level, while up and down increase and decrease the window, respectively.
- Alpha sphere + trigger + joystick: moving up will increase alpha, while moving the joystick down decreases the alpha value.

All the interactions adhere to a consistent pattern: the intersection of the ray with an element, coupled with the direction of the joystick, determines the primary (or more expected) modification to the selected element. Consequently, this pattern yields diverse interactions depending on the specific context. To provide visual feedback, the detected intersection is visually indicated by a dark red color, effectively communicating the interaction to the user. The planes' translations are quite intuitive: they move in the direction of the normal. Planes rotate around an axis that crosses their center in the direction of the viewer, as illustrated in Figure 3.

3.2. Segmentation guidance

In our implementation, we utilize a region growing algorithm to segment regions and place seeds, with supporting planes. While we have found this approach to be effective, it is possible that similar




Figure 4: The segmentation is shown in real-time, and it can be put into context thanks to the medical images.

controls may be suitable for other algorithms. The features required for this process include: *i*) Start/stop region growing, *ii*) addition of seeds or stoppers, *iii*) painting a region or voxel, and *iv*) segmentation deletion.

The control of the segmentation is carried out using the right controller, since we expect the user to devote more time here than manipulating the wedge. The previous interactions may be active depending on whether the ray is intersecting the planes or not. The controls are listed here:

- Without intersection, we can toggle between "seeds", "stoppers", "erasing": this is achieved using the application menu.
- The grip control toggles on and off the region growing algorithm.
- Intersection + trigger paints using the currently selected ID of the brush.

4. Analysis

To analyze the suitability of our metaphor, we carried out two different analyses. First, we evaluated the frame rates. Second, we carried out a small user study to analyze the usability of the metaphor.

4.1. Performance

Since high frame rates are required in VR applications, to avoid sickness, we also provide a performance analysis on an Nvidia Geforce RTX 3070, with three models, placed at two different distances. The results are shown in Table 1. In all cases, frame rates are either around 90 frames per second, and only when the object is close to the viewer, we observe a frame rate decay. Note that the rendering algorithm we use is a GPU-based ray casting [KW03, EHK*06], which is the standard in volume rendering, and whose bottleneck is at the number of fragments that need to be processed because each fragment generates a new ray that needs to traverse the bounding box of the volumetric model.

4.2. Usability evaluation

We also performed a usability evaluation through a user study, divided into two parts: one for the manipulating the geometry of the

Model	Distance	Only render	Wedge	Wedge + segmentation
spheres	Close	85-89	85-89	85-89
obis	Close	25-30	25-30	25-30
head	Close	40-45	40-45	40-45
spheres	Far	85-89	85-89	85-89
obis	Far	85-89	85-89	85-89
head	Far	85-89	85-89	85-89

Table 1: Performance results with the NVidia 3070. The frame rates are high enough when the model is far, and for closer views, the frame rate decays, although stays above 40 fps for the head model and in the range 25-30 for the obis model.



Figure 5: Manipulation task from the user study. The users had to modify the wedge to align the planes as in the reference image.

wedge tool, and another where the users had to use the tool to edit the medical model. The first test consisted of 4 tasks. In all the tasks, the participants need to modify the geometry of the wedge so that the three planes are aligned as asked. An example of an image given to the users is shown in Figure 5. This implies both rotation and translation of, possibly, all planes. We measured the time of the three first tasks, while the last one was not measured, but had a time limit. The idea was to assess whether the participants can perform a certain task in a limited time. The second test consisted of 8 different tasks. These included two different kinds of editions: changing the alpha value, and modifying the window and the level (as in Figure 6). Like in the previous case, each part had 4 subtasks, and the first three are measured, while for the last one, the time was limited. In both tests, we compared the performance working with the previous version of the wedge tool, and the current one. The order in which users were exposed to the different versions was changed for each user, to avoid a learning effect skewing our data. The procedure was the usual one, with a first briefing, filling and signing a consent form, training, task performance, and final questionnaires.

We gathered 8 participants with ages between 20 and 25 years.

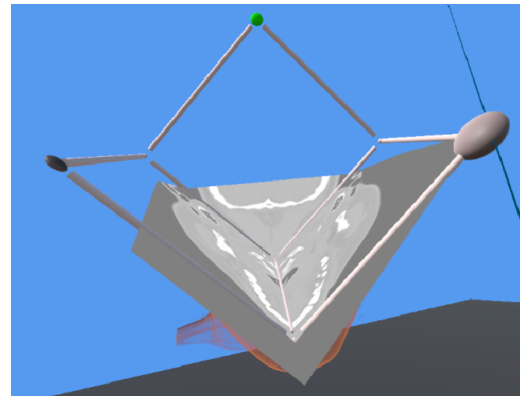


Figure 6: Manipulation example from the study: participants were required to modify the window level to mimic the reference image.

Most of them (5) played video games frequently or every day. The majority (5) had never used a VR setup. None of them had problems carrying out the experiment, and nobody complained of sickness.

The questionnaire regarding the use of the wedge tool had the following questions and average results:

- **Q1:** I have been able to place the wedge as required: 6.57.
- **Q2:** I have found it easy to use the tool to clip the model as required: 6.43.
- **Q3:** I felt unsure at some moments on whether I was performing the task as specified: 1.86.
- **Q4:** I have needed to ask several times how the controls work: 2.43.
- **Q5:** I think that most people would quickly learn the use of the wedge: 6.71.
- **Q6:** It is easy to understand how the tool works: 6.86.
- **Q7:** The tool is easy to use: 6.57
- **Q8:** I could use the tool without help: 6.57.

We also had a last question where we asked the users to grade the experience with the tool between 0 and 10, with a result of 9.

For the Window-Level tool set, the questions, and results were:

- **Q1:** I have been able to modify the model as asked: 6.43.
- **Q2:** I have found it easy to use the tool to modify the model as asked: 6.14.
- **Q3:** I felt unsure at some moments on whether I was performing the task as specified: 2.23.
- **Q4:** I have needed to ask several times how the controls work: 2.57.
- **Q5:** I think that most people would quickly learn the use of the wedge: 6.
- **Q6:** It is easy to understand how the tool works: 6.29.
- **Q7:** The tool is easy to use: 6.57.
- **Q8:** I could use the tool without help: 6.57.

Like in the previous case, we asked the users to grade their experience with the tool between 0 and 10, with a result of 8.86.

The questionnaires show that users felt they could properly achieve the tasks (Q1 and Q2). Moreover, they did not feel unsure

on how they were progressing, and did not require to be reminded of the controls' usage (Q3 and Q4). And the usability questions regarding learnability and ease of usage were also graded quite high.

For the wedge manipulation test, the participants required times varying from less than 10 seconds to around 80 in an outlier case, with averages around 35 for the first task, 19 and 20 for the second and third, respectively. However, the time-limited task was solved in less than 10 seconds by all the users (the time limit was 15 seconds). The times reduced slightly across the tasks, which might be due to a learning effect. We did not perform a statistical analysis since there were only eight users. The second test had a similar behavior, the time averages were around 20, 12, 11, and 9 seconds for the alpha value tasks, and 26, 16, 17, and 14 for the Window Level tasks. Like in the previous case, the time limits (15 and 30, for tasks 4 and 8, respectively) were enough for the participants.

The results strongly suggest that the tool was perceived as useful and usable. During this experiment, we also compared with a previous version of the tool, and the results indicate that the current version, which is simpler and more intuitive, is faster and received superior evaluations by the users. However, the experiment was limited to only 8 users, and their profiles not so varied. As a result, we believe that a second, larger study could be implemented, with a larger number of participants.

5. Conclusions and Future Work

In this paper, we have introduced a novel interaction metaphor specifically designed to facilitate the definition of seed points and guide medical image segmentation algorithms. The tool has been developed for implementation within a virtual reality environment. Its wedge-shaped form enables users to visualize the interior of the volume while maintaining the contextual information provided by three distinct planes. To enhance its versatility, we have incorporated a collection of interaction features that enable effortless modification of the planes and the addition of medical data, accompanied by supplementary tools for adjusting the alpha value, window, and level. To assess its usability, a small-scale user study was conducted, yielding results that indicate a high level of acceptance among participants. Moving forward, our future endeavors involve conducting a second study with a larger participant pool and encompassing a broader range of high-level tasks.

Acknowledgments

Research funded by PID2021-122136OB-C21 from the Ministerio de Ciencia e Innovación, Spain, by 839 FEDER (EU) funds.

References

- [BPC19] BALOUP M., PIETRZAK T., CASIEZ G.: Raycursor: A 3d pointing facilitation technique based on raycasting. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), pp. 1–12. doi:<https://doi.org/10.1145/3290605.3300331>. 2
- [EHK*06] ENGEL K., HADWIGER M., KNISS J., REZK-SALAMA C., WEISKOPF D.: *Real-Time Volume Graphics*. CRC Press, 2006. 3
- [FP00] FRÖHLICH B., PLATE J.: The cubic mouse: a new device for three-dimensional input. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2000), pp. 526–531. 2
- [KBB11] KOPPER R., BACIM F., BOWMAN D. A.: Rapid and accurate 3d selection by progressive refinement. In *2011 IEEE symposium on 3D user interfaces (3DUI)* (2011), IEEE, pp. 67–74. doi:[10.1109/3DUI.2011.5759219](https://doi.org/10.1109/3DUI.2011.5759219). 2
- [KW03] KRUGER J., WESTERMANN R.: Acceleration techniques for gpu-based volume rendering. In *IEEE Visualization, 2003. VIS 2003.* (2003), IEEE, pp. 287–292. 3
- [LNBK*16] LATIFI-NAVID M., BILEN M., KONUKSEVEN E. I., DOĞAN M., ALTUN A.: Fast and accurate semiautomatic haptic segmentation of brain tumor in 3d mri images. *Turkish Journal of Electrical Engineering and Computer Sciences* 24, 3 (2016), 1397–1411. 1
- [LSC*21] LUNGU A. J., SWINKELS W., CLAESSEN L., TU P., EGGER J., CHEN X.: A review on the applications of virtual reality, augmented reality and mixed reality in surgical simulation: an extension to different kinds of surgery. *Expert review of medical devices* 18, 1 (2021), 47–62. 1
- [LYS20] LU Y., YU C., SHI Y.: Investigating bubble mechanism for ray-casting to improve 3d target acquisition in virtual reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2020), IEEE, pp. 35–43. doi:[10.1109/VR46266.2020.00021](https://doi.org/10.1109/VR46266.2020.00021). 2
- [MHG*23] MASLYCH M., HMAITI Y., GHAMANDI R., LEBER P., KATTOJU R. K., BELGA J., LAVIOLA J. J.: Toward intuitive acquisition of occluded vr objects through an interactive disocclusion mini-map. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)* (2023), pp. 460–470. doi:[10.1109/VR55154.2023.00061](https://doi.org/10.1109/VR55154.2023.00061). 2
- [MV23] MOLINA E., VÁZQUEZ P.-P.: Two-step techniques for accurate selection of small elements in vr environments. *Graphical Models* 128 (2023), 101183. 2
- [MVN13] MONCLÚS E., VÁZQUEZ P.-P., NAVAZO I.: DAAPMed: A data-aware anchor point selection tool for medical models in vr environments. In *International Conference on Computer Graphics Theory and Applications* (2013), vol. 2, SCITEPRESS, pp. 308–317. 2
- [OMN*23] ORELLANA B., MONCLÚS E., NAVAZO I., BENDEZÚ Á., MALAGELADA C., AZPIROZ F.: End to end colonic content assessment: Colonometry application. *Diagnostics* 13, 5 (2023). doi:[10.3390/diagnostics13050910](https://doi.org/10.3390/diagnostics13050910). 2
- [PHK04] PIEPER S., HALLE M., KIKINIS R.: 3d slicer. In *2004 2nd IEEE international symposium on biomedical imaging: nano to macro (IEEE Cat No. 04EX821)* (2004), IEEE, pp. 632–635. 1
- [SHT19] SOMPAGNIMDI M. T., HURTER C., TELEA A. C.: Interactive obstruction-free lensing for volumetric data visualization. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019). 2
- [SKM*23] SCHÖN D., KOSCH T., MÜLLER F., SCHMITZ M., GÜNTHER S., BOMMHAARDT L., MÜHLHÄUSER M.: Tailor twist: Assessing rotational mid-air interactions for augmented reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023), pp. 1–14. 2
- [TGK*17] TOMINSKI C., GLADISCH S., KISTER U., DACHSELT R., SCHUMANN H.: Interactive lenses for visualization: An extended survey. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 173–200. 2
- [THY*19] TU H., HUANG S., YUAN J., REN X., TIAN F.: Crossing-based selection with virtual reality head-mounted displays. In *Proceedings of the 2019 CHI conference on human factors in computing systems* (2019), pp. 1–14. 2
- [TPL*22] TRIMPL M. J., PRIMAKOV S., LAMBIN P., STRIDE E., VALIS K. A., GOODING M.: Beyond automatic medical image segmentation – the spectrum between fully manual and fully automatic delineation. *Physics in Medicine and Biology* (2022). doi:[10.1088/1361-6560/ac6d9c](https://doi.org/10.1088/1361-6560/ac6d9c). 2
- [WLC*22] WANG R., LEI T., CUI R., ZHANG B., MENG H., NANDI A. K.: Medical image segmentation using deep learning: A survey. *IET Image Processing* 16, 5 (2022), 1243–1267. 1