POLYTECHNIC OF TURIN

Master's Degree in Biomedical Engineering

Master's Degree Thesis

Mixed Reality Surgical Guidance for Needle Insertion in Thoracic Biopsy Using Meta Quest 3

Supervisors

Prof. Massimo SALVI

Prof. Alberto AREZZO

Eng. Francesco MARZOLA

Dr. Federica BARONTINI

Eng. Matteo PESCIO

Candidate

ALESSIO SCILLIA

July 2026

Declaration

I hereby declare that the content and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

This thesis is licensed under a Creative Commons License, Attribution - Noncommercial - NoDerivative works 4.0 International: see www.creativecommons.org. The text may be reproduced for non-commercial purposes, provided that credit is given to the original author.

Turin, 20/07/2026

AI Use Disclosure

In the writing process of this thesis the following tools have been used to improve the readability and language:

- ChatGPT (OpenAI)
- Google AI Studio (Google)

Turin, 20/07/2026

Abstract

Percutaneous lung biopsy is a key diagnostic procedure for characterizing suspicious pulmonary nodules. However, CT-guided needle insertion is still mainly planned and performed using two-dimensional CT images displayed on external monitors, requiring the interventional radiologist to translate axial image information into a three-dimensional understanding of the lesion, surrounding anatomy, and planned needle trajectory. Augmented Reality (AR) can support this workflow by overlaying patient-specific three-dimensional anatomical models onto the operative field, improving spatial awareness and supporting more intuitive trajectory planning.

This thesis presented the development of a Mixed Reality (MR) guidance system for CT-guided lung biopsy using Meta Quest 3. The work extended a previous Microsoft HoloLens prototype and transferred the application to a more robust mixed reality platform, while improving planning, visualization, registration, and instrument guidance. The system integrated a pre-operative planning environment in 3D Slicer with a Unity application running on Meta Quest 3. In 3D Slicer, a custom module was developed to centralize anatomical segmentation, target definition, trajectory planning, and model export, reducing manual operations and streamlining the workflow from CT data to MR visualization.

In the Unity application, the patient-specific thoracic model was visualized as a hologram and explored by selectively enabling or disabling lungs, bones, vessels, airways, and skin. The pulmonary nodule was kept constantly visible and highlighted in fluorescent green. Automatic registration was implemented using four QR codes positioned around the patient, using CT scanner isocenter as the geometric reference corresponding to the markers center. Real-time CT slice visualization was also integrated through OpenIGTLink to correlate the hologram with the corresponding anatomical section.

The application further included interactive trajectory planning directly in MR. A virtual laser was used to provide visual feedback on whether a candidate needle path reached the pulmonary nodule without intersecting critical anatomical structures, while also estimating the needle insertion depth. In addition, real-time tracking of a cylindrical needle holder was implemented using a binary marker pattern and pose estimation algorithm, providing quantitative feedback on lateral and angular deviation between the planned trajectory and the actual instrument orientation.

The system was evaluated through a phantom-based test and a usability questionnaire completed by two potential clinical users. Marker-based alignment was successfully initialized on the custom thoracic phantom, although the first automatic registration showed an offset that required manual refinement using a

calibration volume. After this correction, the virtual model could be satisfactorily aligned with the physical phantom, confirming the technical feasibility of the QR-code-based registration workflow in a controlled setup. The questionnaire, based on a 5-point Likert scale, showed high perceived usefulness for anatomical visualization and planning: spatial understanding of the target nodule relative to surrounding structures scored 5.00/5, laser-based trajectory guidance scored 4.50/5, and the overall promise of the system for supporting lung biopsy planning and guidance scored 5.00/5. The standardized usability assessment showed a SUS score of 67.5/100, close to the conventional acceptability threshold, while NASA-TLX indicated a moderate workload. Although further validation is required before clinical use, these preliminary results demonstrated the feasibility of implementing an integrated lung biopsy planning and guidance workflow on Meta Quest 3, combining patient-specific MR visualization, marker-based registration, interactive path planning, CT slice visualization, and preliminary needle alignment feedback.

Acknowledgements

I would like to begin by expressing my thanks to those who made this thesis possible, contributing decisively to my professional and personal growth.

My first thanks go to my supervisor, Prof. Salvi. I thank him for the trust he placed in my work, and for his professionalism throughout the entire process.

I would like to mention the Polytechnic of Turin, the institution that has trained me. My university experience has been an intense and stimulating journey, providing me with the foundations and tools to face future challenges with competence and curiosity.

A special dedication goes to the MITIC laboratory, which welcomed me and has been an important environment for my growth. In this context, a particular thank you goes to Prof. Arezzo. I thank him not only for the opportunity to develop my thesis in a research center, but above all for being a source of inspiration. His vision, passion, and ability to motivate are contagious and allowed me to deeply appreciate the value of our work. Thanks to him, I had the privilege of coming into direct contact with the hospital environment, working in a unique interdisciplinary setting surrounded by surgeons and engineers.

A special thanks also go to Eng. Marzola, who closely guided me in every technical aspect of this project. Thank you for your patience during difficult times, for the continuous exchange of ideas, and for your constant availability. You have been an indispensable point of reference.

Table of Contents

List of Tables x

List of Figures xi

Acronyms xv

1 Introduction 1
- 1.1 Context and problem definition . . . 1
- 1.2 Motivation of the thesis . . . 2
- 1.3 Aim of the thesis . . . 2

2 Background 5
- 2.1 Percutaneous Lung Biopsy . . . 5
  - 2.1.1 CT-Guided Procedural Workflow . . . 5
  - 2.1.2 Radiation exposure and dose optimization . . . 8
  - 2.1.3 Procedural complications and technical limitations . . . 9
- 2.2 Navigation systems for interventional guidance . . . 10
  - 2.2.1 Robotic and mechanically assisted guidance systems . . . 11
  - 2.2.2 Electromagnetic tracking-based navigation systems . . . 13
  - 2.2.3 Optical and infrared tracking-based navigation systems . . . 14
  - 2.2.4 Common limitations of navigation systems . . . 15
- 2.3 Mixed and augmented reality in percutaneous interventions . . . 15
  - 2.3.1 Head-mounted displays for extended reality . . . 16
  - 2.3.2 Medical image segmentation for patient-specific virtual models 17
  - 2.3.3 Patient-specific 3D visualization, trajectory guidance, and instrument tracking . . . 18
  - 2.3.4 Challenges and limitations of mixed reality in interventional procedures . . . 20
- 2.4 Critical summary and methodological gap . . . 23

3 Materials and Methods 25
3.1 Tools used . . . 25
3.1.1 CT Dataset . . . 25
3.1.2 3D Slicer . . . 26
3.1.3 TotalSegmentator . . . 27
3.1.4 Unity and Meta XR SDK . . . 28
3.1.5 Meta Quest 3 . . . 30
3.1.6 OpenIGTLink . . . 31
3.1.7 Cylindrical marker for instrument tracking . . . 33
3.1.8 Flask server and OpenCV . . . 35
3.2 Pre-operative planning workflow in 3D Slicer . . . 35
3.2.1 CT loading and visualization . . . 36
3.2.2 Automatic segmentation and anatomical model generation . 36
3.2.3 Target definition, obstacle model, and trajectory analysis . . 40
3.2.4 OpenIGTLink configuration for slice synchronization . . . 45
3.2.5 Export for mixed-reality visualization . . . 45
3.3 Mixed-reality application in Unity for Meta Quest 3 . . . 46
3.3.1 Automatic import of the patient-specific 3D model . . . 47
3.3.2 Graphical user interface and anatomical model management 49
3.3.3 QR-based automatic registration . . . 51
3.3.4 CT slice visualization through OpenIGTLink . . . 52
3.3.5 Virtual laser guidance for manual trajectory selection . . . 54
3.3.6 Camera-based tracking of the real needle . . . 56
3.4 Experimental evaluation . . . 57
3.4.1 Phantom setup . . . 58
3.4.2 Clinical user questionnaire . . . 60

4 Results 62
4.1 Phantom-based evaluation of QR-based alignment . . . 62
4.2 Clinical user questionnaire results . . . 64
4.2.1 Custom questionnaire results . . . 64
4.2.2 System Usability Scale results . . . 66
4.2.3 NASA-TLX results . . . 66

5 Discussion 68
5.1 Interpretation of the preliminary results . . . 68
5.1.1 Interpretation of the phantom-based alignment test . . . 68
5.1.2 Interpretation of questionnaire scores . . . 69
5.2 Main limitations of the study . . . 71
5.2.1 Restricted Validation and Lack of Clinical Evidence . . . 71
5.2.2 Limited Anatomical Realism of the Phantom Model . . . 71

5.2.3 Registration accuracy and manual alignment refinement . . . 72
5.2.4 Limited Accuracy and Robustness of Needle Tracking 72
5.2.5 Computational Constraints of Standalone Implementation . 73
5.2.6 Lack of Validation in Real Clinical Environment 73
5.3 Future developments 73

6 Conclusions 77

Bibliography 79

List of Tables

2.1 Main dosimetric quantities used to describe radiation exposure in Computed Tomography (CT)-guided procedures. 8
2.2 Comparison of representative navigation approaches for image-guided percutaneous interventions. 11
2.3 Representative examples of patient-specific visualization, trajectory guidance, and instrument tracking approaches in Augmented Reality (AR)- and Mixed Reality (MR)-assisted procedures. 21

3.1 Main technical specifications of Meta Quest 3 relevant to the proposed mixed-reality workflow [52, 53, 54, 55]. 32

4.1 Clinical user questionnaire results regarding system features and clinical applicability. Scores are reported as mean values ± standard deviation on a 5-point Likert scale. Negatively worded items are marked with an asterisk (*). 65
4.2 Clinical user questionnaire results regarding system comfort and usability. Scores are reported as mean values ± standard deviation on a 5-point Likert scale. Negatively worded items are marked with an asterisk (*). 66
4.3 NASA-TLX workload results. The final score corresponds to the mean NASA-TLX score. 66

List of Figures

1.1 Conceptual rationale of the thesis. CT images are used to reconstruct a patient-specific three-dimensional anatomical model, visualize anatomical information in mixed reality, and define a planned needle trajectory in relation to the patient or procedural phantom. 3

2.1 Schematic representation of the conventional CT-guided Percutaneous Transthoracic Lung Biopsy (PTLB) workflow. The procedure includes pre-procedural assessment, trajectory planning, patient positioning, skin entry point selection, needle or introducer advancement, repeated limited CT acquisitions for trajectory verification, tissue sampling, and post-procedural imaging. 7

2.2 Robot-assisted navigation system for CT-guided percutaneous lung procedures. The robotic arm provides a physical guide for needle insertion along the planned trajectory. Reproduced from Malik et al. [14]. 12

2.3 Electromagnetic navigation interface of the IMACTIS CT-Navigation system. The workstation displays the projected needle trajectory on two perpendicular reconstructed CT views, allowing the operator to assess the planned insertion angle before needle advancement. Reproduced from Lanouzière et al. [18]. 13

2.4 Schematic comparison between Optical See-Through (OST) and Video See-Through (VST) head-mounted displays. 17

2.5 Schematic representation of trajectory guidance parameters. The planned trajectory defines the intended path from the skin entry point to the target, whereas the actual needle axis may deviate from this path during insertion. Relevant quantities include lateral error, angular error, and distance from the target. 19

3.1 Examples of lesions satisfying the Lung Image Database Consortium and Image Database Resource Initiative (LIDC-IDRI) definitions of (a) nodule $\geq$ 3 mm, (b) nodule $<$ 3 mm, and (c) non-nodule $\geq$ 3 mm [40]. 25
3.2 Standard 3D Slicer views for thoracic CT volume analysis: 3D view with slicing planes (left) and 2D axial, coronal, and sagittal views (right). 27
3.3 Meta XR Simulator running in Unity Editor, simulating the Meta Quest 3 environment. 29
3.4 Meta Quest 3 headset. 30
3.5 Preliminary prototype of the cylindrical visual marker inspired by Cartucho et al. [60]. In this early implementation, the marker pattern was applied to a pen used as a cylindrical surrogate of the biopsy instrument. 34
3.6 Pre-operative planning workflow implemented in the custom 3D Slicer module. 37
3.7 Custom 3D Slicer module interface for pre-operative planning. The interface centralizes the main operations required for thoracic biopsy planning, including CT loading, segmentation, target definition, trajectory analysis, and scene export for mixed-reality visualization. 38
3.8 Visualization in 3D Slicer of the segmentation outputs generated with the four TotalSegmentator tasks used in the planning workflow: (a) main thoracic anatomical structures obtained with the `total` task; (b) candidate pulmonary nodule obtained with the `lung_nodules` task; (c) pulmonary vessels and airways obtained with the `lung_vessels` task; and (d) external body surface obtained with the `body` task. 39
3.9 Manual definition of the target nodule in 3D Slicer using the Markups system. A fiducial point is placed at the desired biopsy target location on the CT slices. 41
3.10 List of the anatomical structures available for selection (left) and an example of the resulting obstacle model (right) generated by combining the segmented ribs, pulmonary vessels, and airways. This model was used during trajectory analysis to identify feasible paths from the skin surface to the target nodule. 41
3.11 Percutaneous Approach Analysis interface integrated into the custom planning workflow. The parameter section allows the user to select the target point, skin surface, and obstacle model, while the outcome section displays the accessibility map, candidate paths, shortest and longest trajectories, and selected trajectory information. 42

3.12 Colorimetric accessibility map generated on the skin surface. Redder regions indicate shorter accessible paths to the target, whereas lighter regions correspond to longer paths. 44
3.13 Visual verification of the planned trajectory in 3D Slicer. The virtual instrument is oriented along the selected path from the skin entry point to the target nodule, allowing a visual check of the planned trajectory. 44
3.14 Examples of the patient-specific anatomical models visualized in mixed reality in Unity on Meta Quest 3. Different anatomical configurations could be displayed according to the visualization needs, including airways and vessels (AWVessels), bones, lungs with transparent skin, and the complete thoracic model. 48
3.15 Graphical user interface developed in Unity for the Meta Quest 3 application. The interface provides controls for OpenIGTLink-based slice visualization, model manipulation, QR-based registration, manual trajectory planning, anatomical structure visibility, opacity regulation, and needle tracking. 49
3.16 Example of QR-based automatic registration. The patient-specific anatomical model is aligned in MR with the four-QR-code reference system. 52
3.17 Dynamic CT slice visualization in the mixed-reality application. A movable slicing plane intersects the imported anatomical model, while the corresponding CT section is displayed on a virtual screen next to the graphical user interface. 53
3.18 Virtual laser guidance for manual trajectory selection in the mixed-reality application. (a) General view of the grabbable sphere used to control the candidate trajectory. (b) Valid trajectory reaching the target nodule, shown in green. (c) Invalid trajectory intersecting the ribs, shown in red. (d) Skin entry point displayed after trajectory confirmation. 55
3.19 Camera-based needle tracking and preliminary guidance feedback. The tracked needle axis is compared with the manually planned trajectory, allowing the estimation of lateral deviation and angular deviation. 57
3.20 Custom thoracic phantom used for the preliminary evaluation. . . . 58
3.21 Phantom setup used for the preliminary experimental evaluation. The phantom consisted of a gelatin matrix containing artificial nodules and 3D-printed bones. A LEGO calibration volume and a four-QR-code reference pattern were included to support mixed-reality registration and manual alignment refinement. 59

4.1 Mixed-reality visualization before and after manual alignment refinement using the calibration volume as a spatial reference. 63

Acronyms

CT Computed Tomography

AR Augmented Reality

MR Mixed Reality

VR Virtual Reality

XR Extended Reality

PTLB Percutaneous Transthoracic Lung Biopsy

US Ultrasound

PET Positron Emission Tomography

FNA Fine Needle Aspiration

CNB Core Needle Biopsy

$\mathbf{CTDI_{vol}}$ Computed Tomography Dose Index Volume

DLP Dose Length Product

SSDE Size-Specific Dose Estimate

ALARA As Low As Reasonably Achievable

TTNA Transthoracic Needle Aspiration

PIGA Precise Intelligent Guide Arm

HMD Head-Mounted Display

MRI Magnetic Resonance Imaging

OST Optical See-Through

VST Video See-Through

LIDC-IDRI Lung Image Database Consortium and Image Database Resource Initiative

DIEP Deep Inferior Epigastric Perforator

CTA Computed Tomography Angiography

MRUK Mixed Reality Utility Kit

IPD Inter-Pupillary Distance

API Application Programming Interface

OpenIGTLink Open Image-Guided Therapy Link

NIfTI Neuroimaging Informatics Technology Initiative

SUS System Usability Scale

NASA-TLX NASA Task Load Index

Chapter 1

Introduction

1.1 Context and problem definition

CT-guided lung biopsy is a key minimally invasive procedure for the diagnosis and characterization of suspicious pulmonary nodules. During the procedure, the interventional radiologist plans a safe percutaneous access route on CT images and manually advances the needle toward the target lesion through the chest wall. Although this technique is clinically established, it remains technically demanding because the operator must translate two-dimensional image information into a three-dimensional understanding of the patient anatomy and of the planned needle trajectory.

This spatial translation is particularly challenging when the target nodule is small, deep, poorly visible, or located close to critical anatomical structures such as ribs, vessels, airways, or the pleural surface. In conventional CT-guided workflows, needle advancement is usually verified through repeated image acquisitions rather than through continuous real-time visualization. As a result, deviations from the planned trajectory may require additional corrections and control scans, increasing procedural complexity, time, and radiation exposure. Respiratory motion and tissue deformation further contribute to the difficulty of maintaining correspondence between the pre-procedural image information and the actual anatomical configuration during the procedure.

These limitations highlight the need for guidance strategies capable of improving spatial understanding, supporting trajectory planning, and helping the operator relate the target lesion, surrounding anatomical structures, skin entry point, and needle path within the same three-dimensional reference frame.

1.2 Motivation of the thesis

Navigation and image-guidance systems have been proposed to support percutaneous procedures by improving the relationship between medical images, patient anatomy, and the interventional instrument. However, many existing solutions still present guidance information on external workstations or monitors, requiring the operator to alternate attention between the patient, the needle, and the navigation display. This does not fully eliminate the cognitive and ergonomic gap between image-based planning and manual action in the procedural field.

MR offers a possible way to reduce this gap by displaying patient-specific anatomical information directly within the operator's field of view. In the context of CT-guided lung biopsy, a mixed-reality system could allow the operator to visualize the target nodule, surrounding anatomical structures, and planned trajectory as interactive three-dimensional content related to the physical patient or to a procedural phantom. This could support a more intuitive understanding of the access route and provide an additional spatial reference during planning and guidance.

From a technological perspective, the development of such a system requires the integration of multiple components into a coherent workflow. These include medical image processing, anatomical segmentation, trajectory planning, export of patient-specific models, registration between virtual and physical spaces, interactive visualization, and preliminary instrument tracking. Therefore, the motivation of this thesis is not limited to the visualization of a three-dimensional model, but concerns the development of an integrated workflow connecting pre-operative planning and mixed-reality guidance.

A further motivation is related to the hardware platform. Many previous medical AR and MR prototypes have been developed for Microsoft HoloLens-based systems. In this thesis, the workflow was transferred and extended to Meta Quest 3, a video pass-through mixed-reality headset that provides color passthrough, Unity compatibility, standalone execution, and accessible hardware for research prototyping. Although Meta Quest 3 is not a medical device, it represents a suitable platform for investigating the technical feasibility of mixed-reality guidance workflows in controlled experimental settings.

1.3 Aim of the thesis

The aim of this thesis is to develop and preliminarily evaluate a Meta Quest 3-based MR system for patient-specific planning and guidance of CT-guided thoracic needle insertion. The proposed workflow integrates a pre-operative planning environment in 3D Slicer with a Unity application running on Meta Quest 3. An overview of

the conceptual workflow is shown in Figure 1.1.

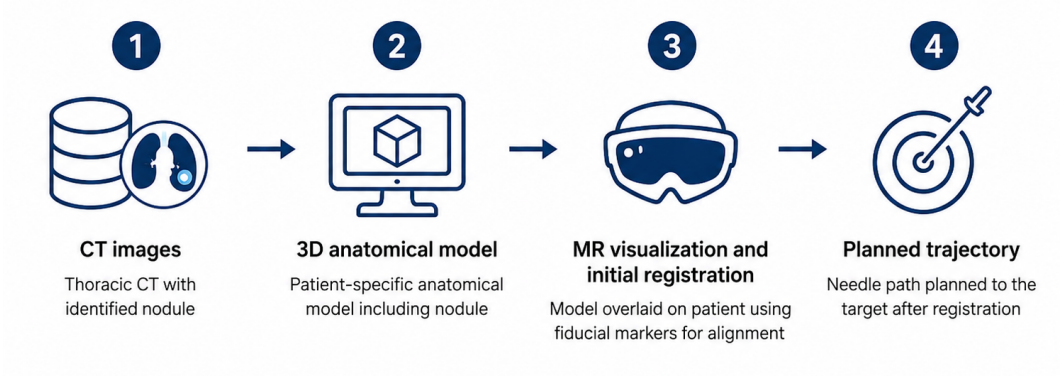

Figure 1.1: Conceptual rationale of the thesis. CT images are used to reconstruct a patient-specific three-dimensional anatomical model, visualize anatomical information in mixed reality, and define a planned needle trajectory in relation to the patient or procedural phantom.

The specific objectives of the thesis are:

1. to develop a dedicated 3D Slicer workflow for CT loading, anatomical segmentation, target definition, trajectory planning, and export of patient-specific thoracic models;
2. to implement a Unity-based mixed-reality application for Meta Quest 3 capable of importing and visualizing the exported anatomical models;
3. to provide interactive visualization of clinically relevant structures directly in the mixed-reality environment, including lungs, target nodule, skin surface, bones, vessels, airways, and planned trajectory;
4. to implement a QR-code-based registration strategy for initial alignment between the virtual model and the patient or procedural reference;
5. to integrate dynamic CT slice visualization through OpenIGTLink, allowing sectional image information to be inspected in relation to the three-dimensional model;
6. to implement interactive mixed-reality tools for manual trajectory planning with a laser-based system, and preliminary needle tracking feedback;
7. to preliminarily evaluate the technical feasibility of the proposed workflow through a phantom-based test and a clinical user questionnaire.

The goal of this work is not to provide clinical validation of mixed-reality-guided lung biopsy, but to demonstrate the feasibility of an integrated Meta Quest 3-based platform combining patient-specific visualization, registration, trajectory planning, and preliminary guidance support for future MR-assisted thoracic interventions.

Chapter 2

Background

2.1 Percutaneous Lung Biopsy

PTLB is a minimally invasive technique for obtaining tissue samples from suspicious peripheral pulmonary lesions using a needle passing percutaneously through the chest wall, generally through an intercostal space. Its main clinical role is to determine whether a suspected lesion is benign or malignant, classify tumor subtype, and obtain adequate material for ancillary analyses, including immunohistochemical and molecular testing [1].

Among the available imaging modalities, CT is the most commonly used guidance technique for percutaneous lung biopsy and represents the imaging reference for this thesis. CT provides cross-sectional visualization of the thorax, allowing accurate localization of the target lesion and assessment of its spatial relationship with the surrounding anatomical structures. This makes CT particularly suitable for deep pulmonary nodules or lesions that are not visible with ultrasound. Ultrasound (US) guidance may be used for selected lesions abutting the pleural surface, but its applicability is limited for deeper intrapulmonary targets because aerated lung tissue strongly attenuates ultrasound propagation [2].

2.1.1 CT-Guided Procedural Workflow

In a conventional CT-guided PTLB workflow, the procedure begins with a pre-procedural assessment aimed at confirming the indication for biopsy and minimizing procedural risks. Coagulation parameters are checked, and previously acquired CT or PET/CT images are reviewed to evaluate the target lesion, its relationship with surrounding anatomical structures, and the feasibility of a safe percutaneous approach [1, 3]. CT imaging plays a central role in this phase because it allows the operator to select the most appropriate target, define the safest access route, and plan the needle trajectory while avoiding structures such as ribs, vessels, bronchi,

fissures, and bullae. When available, PET/CT may provide additional functional information by identifying the metabolically active portion of heterogeneous or poorly defined lesions [1, 4].

After the target and access route have been selected, the patient is positioned on the CT table according to lesion location, planned trajectory, and procedural feasibility. A first CT acquisition is then performed to confirm the lesion location and identify the axial level for needle insertion. The skin entry point is selected by correlating the CT images with physical references on the patient's body. In conventional practice, this can be achieved using the gantry laser and radiopaque skin markers or grids, which allow the operator to transfer the planned entry point from the image slice to the patient surface [1].

Following sterile preparation and local anesthesia, the needle or introducer is advanced from the marked skin entry point toward the target. Tissue sampling can be performed using either Fine Needle Aspiration (FNA) or Core Needle Biopsy (CNB), depending on the diagnostic question and the amount of tissue required. CNB is often preferred when histological characterization, immunohistochemical analysis, or molecular testing is needed even though is more invasive than FNA, and it is frequently performed using a coaxial technique to obtain multiple samples through a single pleural access [1, 3, 5].

In conventional CT guidance, needle progression is not continuously visualized. Instead, the operator advances the device step by step and acquires repeated limited CT scans to verify its depth, orientation, and relationship with the target. If the needle does not follow the planned trajectory, its direction must be corrected before further advancement. This iterative acquisition-and-correction workflow is a key limitation of conventional CT-guided biopsy, because accurate needle placement depends on the operator's ability to translate the planned trajectory from two-dimensional image slices to the patient and to adjust the needle based on intermittent imaging feedback [1, 3]. CT fluoroscopy can provide more continuous visualization of needle progression, but at the cost of increased radiation exposure to both the patient and the operator [5].

At the end of the sampling phase, additional imaging is performed to evaluate the immediate outcome of the procedure and detect early complications, particularly pneumothorax and pulmonary hemorrhage [3, 6]. The main steps of the conventional CT-guided lung biopsy workflow are summarized in Figure 2.1.

Overall, CT guidance provides accurate anatomical information throughout the biopsy workflow, from planning to post-procedural assessment. However, this accuracy is achieved through repeated image acquisitions used to verify needle position and guide trajectory corrections. For this reason, radiation exposure and dose optimization represent central issues in conventional CT-guided lung biopsy.

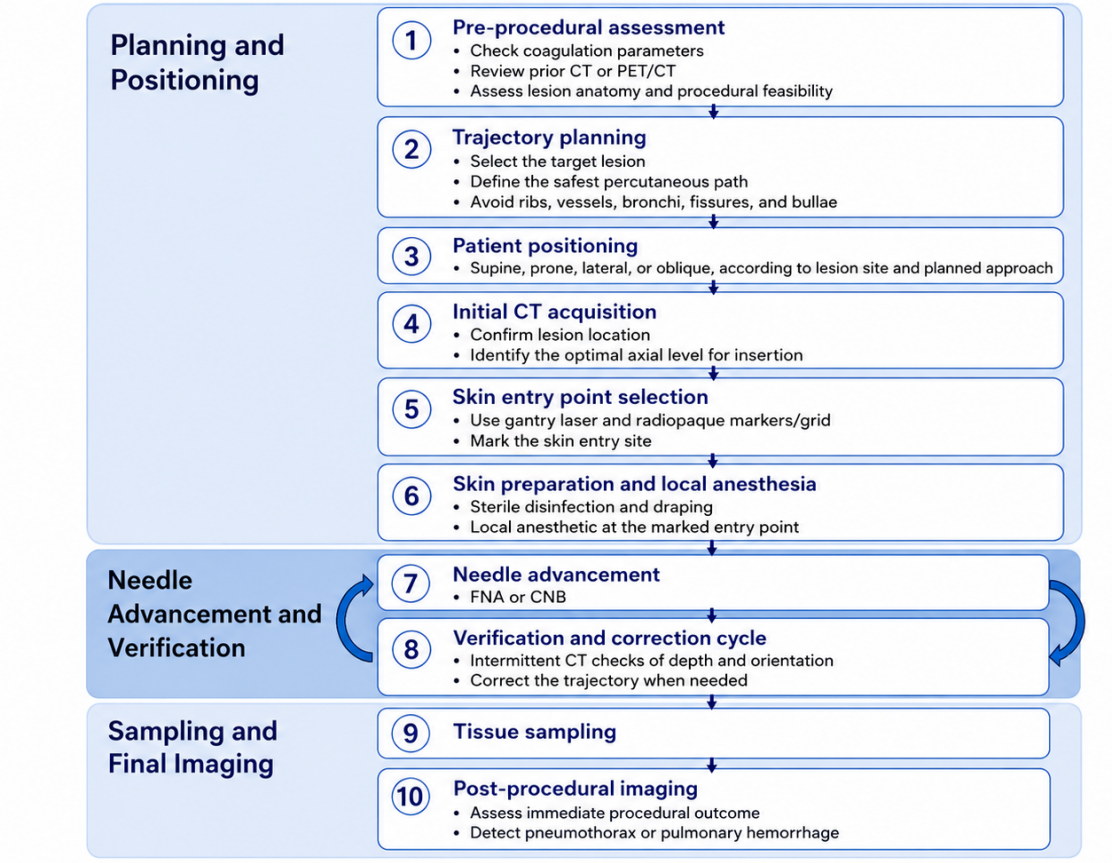

Figure 2.1: Schematic representation of the conventional CT-guided PTLB workflow. The procedure includes pre-procedural assessment, trajectory planning, patient positioning, skin entry point selection, needle or introducer advancement, repeated limited CT acquisitions for trajectory verification, tissue sampling, and post-procedural imaging.

2.1.2 Radiation exposure and dose optimization

One of the main limitations of conventional CT-guided PTLB is the use of ionizing radiation. Radiation exposure is not limited to a single diagnostic acquisition, but results from the combination of the initial planning scan, multiple intra-procedural acquisitions for needle guidance, and the final post-biopsy control scan. Since conventional CT guidance does not provide continuous visualization of needle progression, the operator usually advances the needle step by step and performs additional limited CT acquisitions to confirm its depth, orientation, and relationship with the target lesion. As a result, the total radiation burden depends not only on the scan protocol itself, but also on procedural complexity, lesion size and depth, respiratory motion, and number of needle corrections.

Several dosimetric quantities can be used to describe radiation exposure in CT-guided procedures. The most commonly reported metrics are summarized in Table 2.1.

Table 2.1: Main dosimetric quantities used to describe radiation exposure in CT-guided procedures.

Metric	**Unit**	**Meaning**
$CTDI_{vol}$	mGy	Standardized index of scanner radiation output for a single acquisition.
DLP	mGy·cm	Product of $CTDI_{vol}$ and scan length, reflecting the radiation output over the scanned anatomical region.
SSDE	mGy	Patient-size-adjusted dose estimate, providing a more individualized assessment than $CTDI_{vol}$.
Effective dose	mSv	Risk-weighted estimate of whole-body exposure.

In a retrospective dosimetric analysis of CT-guided biopsy procedures, Siomou et al. highlighted that the dose associated with lung biopsy is divided between more extensive helical scans—generally used for planning or anatomical assessment—and localized scans used during the procedure itself [7]. Although the latter typically involve a smaller scan volume, their repetition as the needle advances can contribute significantly to cumulative exposure. This is particularly relevant in technically complex cases, where multiple checks and corrections are necessary before reaching the target.

CT fluoroscopy has been introduced to provide more continuous or near real-time visualization of needle progression. This can facilitate targeting, reduce the need for repeated needle repositioning, and potentially shorten the procedure. However, this advantage comes at the cost of increased radiation exposure, particularly when fluoroscopic guidance is used extensively. Nakamura et al. reported higher

estimated effective dose for both patients and physicians during CT fluoroscopy-guided lung biopsy compared with conventional CT guidance, despite comparable diagnostic performance and a lower complication rate [5]. Similarly, Weinand et al. reported comparable non-diagnostic sample rates between CT fluoroscopy and conventional CT guidance, but a significantly higher total DLP in the CT fluoroscopy group [6]. Therefore, CT fluoroscopy may improve procedural feedback, but its use requires careful balancing between guidance accuracy and radiation protection.

For this reason, dose optimization is a central aspect of CT-guided lung biopsy and follows the principle of As Low As Reasonably Achievable (ALARA), while maintaining sufficient image quality for safe needle guidance. Optimization strategies include reducing tube voltage and current, limiting scan length to the region of interest, decreasing the number of acquisitions, using intermittent rather than continuous imaging when possible, and adopting low-dose or ultra-low-dose protocols. Kallianos et al. showed that a low-dose CT-guided lung biopsy protocol based on modifications of kVp, mA, and scan coverage substantially reduced the mean DLP without significant changes in diagnostic outcome or complication rate [8]. However, even though low-dose protocols mainly reduce the dose of each acquisition, they do not completely eliminate the need for repeated imaging during the procedure. Consequently, the total radiation exposure remains strongly influenced by the number of scans required to plan, correct, and confirm the needle trajectory. This represents an important limitation of conventional CT-guided workflows and provides a rationale for navigation and mixed reality systems that aim to transfer pre-procedural imaging information into the procedural field, support trajectory planning, and potentially reduce the number of intermediate CT acquisitions needed during needle advancement.

2.1.3 Procedural complications and technical limitations

Although CT-guided PTLB is widely recognized as an accurate and minimally invasive diagnostic procedure, it is not without risks. Most complications are related to the need to cross the pleura and advance the needle through aerated lung parenchyma to reach the target lesion. Therefore, procedural safety depends not only on patient and lesion characteristics, but also on the ability to define and reproduce a safe needle trajectory during the intervention.

The most common complication is pneumothorax, caused by the entry of air into the pleural space after pleural or parenchymal injury. Its incidence varies widely across studies, depending on patient characteristics, lesion features, biopsy technique, and detection criteria. Pneumothorax is often self-limited, but in a smaller proportion of cases chest drainage may be required [2, 5, 6]. Pulmonary hemorrhage and hemoptysis represent the second main category of complications.

They are generally less frequent than pneumothorax and are often limited, but they may become clinically relevant when bleeding occurs along the needle path or into the airways [2, 3]. Rare but potentially severe adverse events, such as systemic air embolism, hemothorax, tumor seeding, or severe respiratory and cardiovascular events, have also been reported [1, 3].

From a procedural point of view, several risk factors are directly related to the geometry and stability of the needle path. A long intrapulmonary trajectory requires the needle to traverse a greater amount of aerated lung tissue, increasing the risk of parenchymal injury and air leakage. Repeated pleural punctures can further increase the probability of pneumothorax, while multiple needle corrections may prolong the procedure and create additional tissue trauma. Similarly, a suboptimal insertion angle can make needle advancement less stable and may require further adjustments to reach the target. Respiratory motion adds an additional source of uncertainty, because the relative position of the lesion, surrounding structures, and needle can change between successive CT acquisitions [1, 5, 6, 9]. These limitations highlight the technical difficulty of translating a trajectory planned on CT images into a precise manual needle movement on the patient. Between acquisitions, needle progression is not directly visualized, and deviations from the intended path may require trajectory corrections. Each correction can increase procedure time, the number of control acquisitions, and the risk of additional pleural or parenchymal injury.

For this reason, complication prevention is not only a matter of clinical management, but also a problem of procedural guidance. A guidance system should support the reproduction of a short and stable trajectory while limiting unnecessary needle corrections. This requirement provides the rationale for navigation systems in interventional radiology, which aim to improve trajectory planning, needle alignment, and procedural control during image-guided interventions.

2.2 Navigation systems for interventional guidance

Navigation systems for percutaneous interventions have been developed to support the operator in reproducing, on the patient, the access route planned on medical images. Their common objective is to improve the spatial relationship between the target lesion, the planned trajectory, and the interventional instrument, thereby reducing the dependence on purely freehand needle placement. In the context of CT-guided interventions, these systems are particularly relevant because they can support trajectory planning, needle alignment, and procedural feedback while potentially reducing needle corrections, control scans, procedure duration, and

radiation exposure [10, 11, 12].

Different navigation approaches can be classified according to the type of support they provide to the operator. Robotic or mechanically assisted systems provide a physical guide for the planned trajectory. Electromagnetic and optical tracking systems estimate the position and orientation of the instrument and display its projected path with respect to the image data. MR-based guidance systems, instead, aim to present anatomical and trajectory information directly within the procedural environment or the operator's field of view. A comparison of these approaches is reported in Table 2.2.

Table 2.2: Comparison of representative navigation approaches for image-guided percutaneous interventions.

Technology	**Tracking method**	**Main advantage**	**Main limitation**
Robotic or mechanically assisted guidance	Robot or mechanical guide aligned to the planned trajectory	Stable physical alignment of the needle path	Dedicated hardware and limited flexibility during manual adjustments
Electromagnetic tracking	Electromagnetic sensor tracked within a generated field	Instrument tracking without optical visibility constraints	Possible distortion from metallic objects or electromagnetic interference
Optical or infrared tracking	Cameras track optical or infrared markers on the patient and instrument	High tracking accuracy and no magnetic field distortion	Marker occlusion and dependence on camera visibility
MR guidance	Marker-, camera-, or tracking-based registration depending on the implementation	Anatomical and trajectory information can be displayed in the operator's field of view	Accuracy still depends on registration, tracking, motion, and visualization stability

2.2.1 Robotic and mechanically assisted guidance systems

Robotic and mechanically assisted guidance systems support percutaneous needle placement by translating an image-based plan into a physical guide. In a typical workflow, the operator selects the target lesion and the skin entry point on CT images. The system then calculates the planned trajectory, including angulation and insertion depth, and aligns a robotic arm or mechanical guide along that path [13].

The hardware generally includes a robotic arm or mechanical positioning device, a planning workstation, and a sterile needle guide mounted on the end-effector.

The feedback provided to the operator is therefore mainly physical: the system constrains or stabilizes the direction of needle insertion according to the planned trajectory. Importantly, these systems do not autonomously perform the biopsy. The radiologist remains responsible for image interpretation, trajectory approval, needle advancement, tissue sampling, and procedural safety, while the robotic platform assists by maintaining the planned orientation.

A representative example is the Precise Intelligent Guide Arm (PIGA) CT system described by Malik et al. In this system, the planned entry point and target are defined on CT images, the robotic arm is aligned with the calculated trajectory, and the operator manually advances the needle through the guide, as shown in Figure 2.2. Compared with conventional CT guidance, robotic assistance has been associated with shorter procedure duration, fewer check scans, fewer needle adjustments, and improved diagnostic efficacy, without a clear increase in complications [14]. Similar results have been reported for other robotic platforms, such as MAXIO, which follow the same principle of robot-assisted alignment rather than autonomous needle insertion [15, 16].

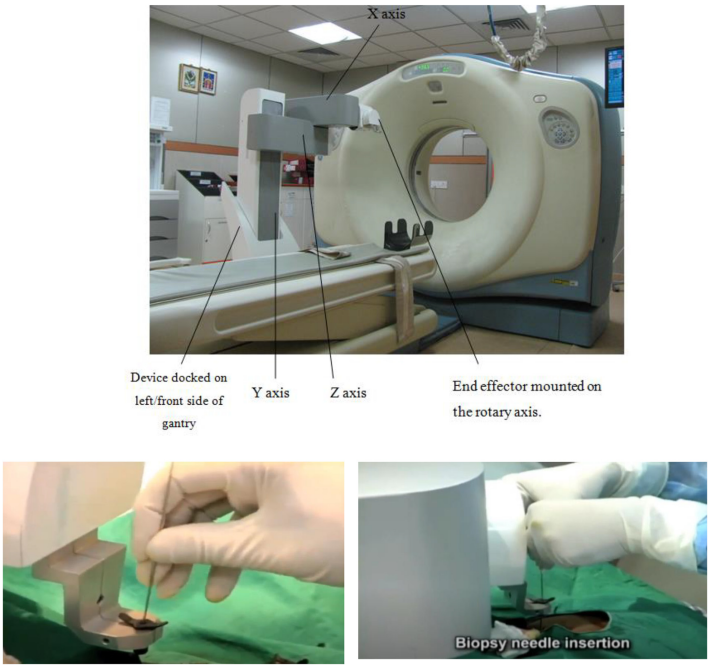

Figure 2.2: Robot-assisted navigation system for CT-guided percutaneous lung procedures. The robotic arm provides a physical guide for needle insertion along the planned trajectory. Reproduced from Malik et al. [14].

The main advantage of robotic guidance is the possibility of stabilizing the planned needle orientation and reducing variability in manual alignment. Its main

limitation is the need for dedicated hardware, room integration, and a workflow that may be less flexible when rapid trajectory adjustments are required.

2.2.2 Electromagnetic tracking-based navigation systems

Electromagnetic tracking-based navigation systems provide guidance by estimating the position and orientation of the instrument within an electromagnetic field. Unlike robotic systems, they do not physically constrain the needle path. Instead, they track the instrument and display its projected trajectory with respect to the CT dataset.

A representative example is the IMACTIS CT-Navigation system. Its essential components are an electromagnetic field generator placed near the patient, a sensor integrated into a needle holder, a navigation workstation, and a registration procedure that links the tracking space to the CT image space. After the initial CT acquisition and registration, the system displays the expected needle path on reconstructed CT views while the operator moves the needle holder on the patient's skin. This allows the radiologist to adjust the entry point and insertion angle before needle advancement, as illustrated in Figure 2.3[17, 18].

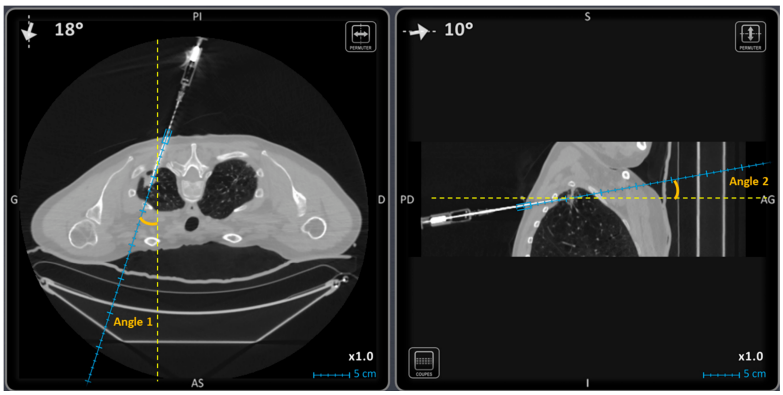

Figure 2.3: Electromagnetic navigation interface of the IMACTIS CT-Navigation system. The workstation displays the projected needle trajectory on two perpendicular reconstructed CT views, allowing the operator to assess the planned insertion angle before needle advancement. Reproduced from Lanouzière et al. [18].

The main advantage of electromagnetic tracking is that it does not require a direct line of sight between cameras and markers. This makes it useful in constrained procedural environments and for access routes where optical visibility may be difficult to maintain. Clinical studies have shown that electromagnetic navigation can improve first needle placement accuracy and reduce the number of additional control CT acquisitions in CT-guided interventions [17]. In percutaneous

lung biopsy, Lanouzière et al. reported similar technical and diagnostic success compared with conventional CT guidance, with significantly lower DLP in the navigation group [18]. The main limitation is the potential distortion of the electromagnetic field caused by metallic objects or electromagnetic interference in the procedural environment.

2.2.3 Optical and infrared tracking-based navigation systems

Optical and infrared tracking-based systems provide navigation by using external cameras to detect markers placed on the patient and on the interventional instrument. As with electromagnetic systems, they do not physically guide the needle. Their feedback is visual: the tracked instrument position and orientation are displayed in relation to the registered CT dataset, allowing the operator to evaluate alignment with the target and planned trajectory during manual needle advancement.

A representative example is the SIRIO system, an optical CT-based navigation platform used for percutaneous lung biopsy [19, 20]. In this workflow, patient tracking is performed using markers placed on the patient's body, while needle tracking is performed through markers or a tracked device associated with the interventional instrument. After an initial CT acquisition, the image dataset is transferred to the navigation workstation and registered with the physical patient space. The system can then display the position and orientation of the needle with respect to the target lesion and the planned trajectory.

In thoracic applications, respiratory motion is a critical issue because lung anatomy and target position change during breathing. For this reason, SIRIO incorporates respiratory phase management by monitoring chest motion and comparing the current respiratory state with the phase acquired during the reference CT scan. This helps the operator perform needle advancement when the anatomical configuration is closer to the one represented in the registered CT model [20, 19].

The main advantage of optical and infrared tracking is that it is not constrained by a physical guide while also avoiding magnetic field distortions and enabling accurate spatial tracking.Clinical studies on SIRIO-assisted lung biopsy have reported encouraging results, including accurate targeting of small pulmonary lesions, reduced procedure time, lower radiation exposure, fewer CT acquisitions, and improved diagnostic success in selected settings [19, 20]. The main limitation is the requirement for a continuous line of sight between the cameras and the markers. Marker occlusion, room layout, and operator positioning can therefore affect tracking reliability.

2.2.4 Common limitations of navigation systems

Although navigation systems can improve the accuracy, efficiency, and safety of CT-guided percutaneous procedures, several limitations still affect their integration into routine clinical workflows [13, 12]. First, most systems require dedicated hardware, such as robotic arms, electromagnetic field generators, infrared cameras, tracking markers, or navigation workstations. These components must be positioned, calibrated, and integrated into the procedural room without interfering with patient access, sterility, or the normal operation of the CT scanner.

A second common limitation is the need for accurate calibration and registration. Navigation systems rely on a stable spatial relationship between the image dataset, the patient, and the interventional instrument. This relationship can be affected by registration errors, patient motion, respiratory deformation, changes in patient position, or soft-tissue displacement during needle advancement. This problem is particularly relevant in thoracic interventions, where respiratory motion can modify the position of the target lesion and surrounding structures between image acquisition and needle insertion [1, 17].

Additional limitations depend on the tracking principle. Optical and infrared systems require continuous marker visibility and can be affected by marker occlusion or camera positioning. Electromagnetic systems avoid line-of-sight constraints, but they can be affected by field distortion or electromagnetic interference. Robotic and mechanically assisted systems provide stable physical alignment, but they introduce additional mechanical components.

Finally, many navigation systems still present guidance information on external workstations or monitors. As a result, the operator must alternate attention between the patient, the needle, and the navigation display. This reduces, but does not completely eliminate, the cognitive and ergonomic gap between image-based planning and manual action on the patient [21, 22]. MR-based guidance addresses this aspect by changing how information is presented to the operator: it can bring patient-specific anatomical models, target information, and trajectory cues into the procedural environment or the operator's field of view. However, MR does not remove the need for accurate registration and reliable tracking. Rather, it depends on these components while providing a different and potentially more intuitive interface for procedural guidance.

2.3 Mixed and augmented reality in percutaneous interventions

Building on the limitations of conventional and navigation-assisted CT-guided procedures, AR and MR can be interpreted as an extension of image-guided

navigation toward a more integrated and spatially intuitive form of procedural support. Rather than presenting guidance information only as numerical values or screen-based views, these technologies aim to display patient-specific anatomical data, target information, and trajectory cues as three-dimensional virtual content that can be explored, manipulated, and, when appropriate, spatially aligned with the procedural scene [21, 23].

In percutaneous interventions, this approach is particularly relevant because procedural success depends on the operator's ability to understand the three-dimensional relationship between the target lesion, surrounding anatomical structures, skin entry point, interventional instrument, and planned access route. For this reason, they represent not only a technological alternative to existing navigation systems, but also a different way of perceiving and using image-based information during needle-based procedures.

2.3.1 Head-mounted displays for extended reality

Head-Mounted Displays (HMDs) are wearable display systems positioned in front of the user's eyes to provide visual information in an immersive or semi-immersive way. They generally include one or more displays, an optical system for image presentation, motion sensors for head tracking, and, depending on the device, cameras or additional sensors for spatial mapping and interaction. In medical AR and MR applications, HMDs are particularly relevant because they allow virtual information to be visualized while the operator maintains awareness of the procedural environment [24, 25].

For AR and MR applications, two main see-through architectures can be distinguished: OST and VST. In OST headsets, such as Microsoft HoloLens and HoloLens 2, the user observes the real world directly through transparent or semi-transparent optical elements, while virtual objects are projected into the field of view through optical combiners or waveguides. This approach preserves a direct view of the environment, but virtual content is optically superimposed onto the real scene and may therefore be affected by limitations in brightness, contrast, field of view, depth perception, and occlusion management [26, 22].

In VST headsets, also referred to as video pass-through devices, the real environment is acquired by external cameras, processed by the device, combined digitally with virtual content, and then displayed on internal screens. Examples of this architecture include devices such as Meta Quest 3 and Apple Vision Pro, as well as video pass-through surgical or research platforms [27, 22]. Since both the real and virtual images are digitally composed before being shown to the user, VST enables greater control over visual blending, opacity, and occlusion between real and virtual elements. However, the perceived quality of the real environment depends on camera resolution, display quality, image processing, tracking stability,

and end-to-end latency [26, 28].

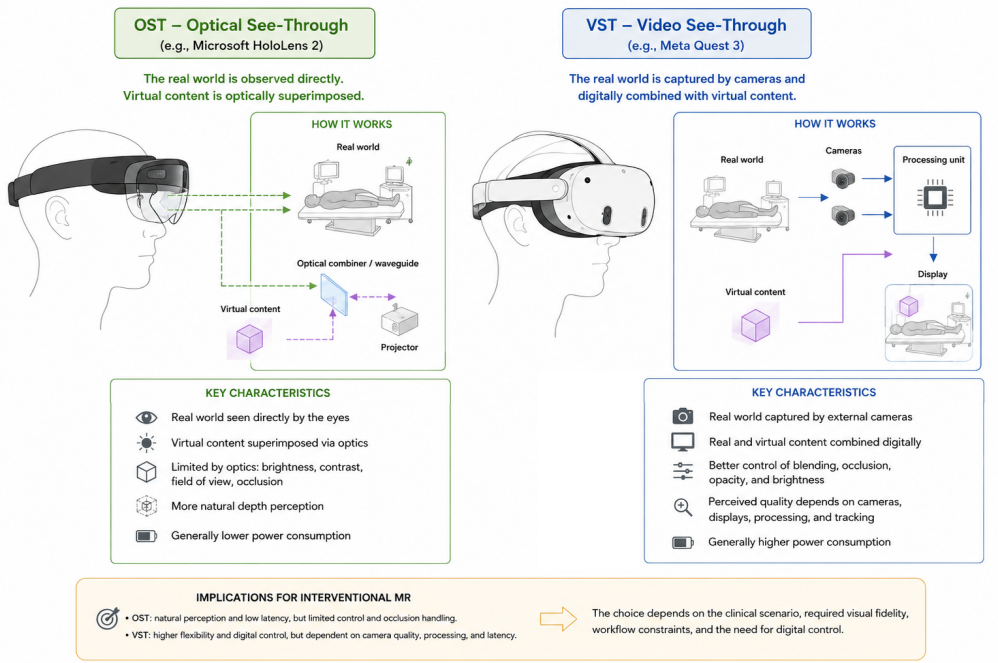

Figure 2.4: Schematic comparison between OST and VST head-mounted displays.

The choice of an HMD therefore depends on several technical factors, including field of view, display resolution, brightness and contrast, tracking accuracy, latency, weight, comfort, computational capabilities, battery life, and supported interaction modalities. These characteristics influence how effectively virtual information can be perceived, manipulated, and spatially related to the real environment. For this reason, the hardware architecture of the HMD represents a fundamental design choice in any AR or MR system, before considering the specific requirements of a clinical or interventional application.

2.3.2 Medical image segmentation for patient-specific virtual models

The use of AR and MR in image-guided procedures requires the conversion of pre-procedural imaging data into virtual information that can be displayed and used during planning or guidance. In the case of CT-based lung biopsy support, the image volume must be processed to identify the target lesion and the anatomical structures that are relevant for trajectory planning, such as the lungs, ribs, vessels, airways, pleural surface, and skin surface. These structures are useful because they help define the safest access route, identify the skin entry point, avoid critical

anatomy, and estimate the spatial relationship between the planned trajectory and the target.

Automated segmentation methods based on deep learning can reduce processing time and improve workflow reproducibility. Tools such as TotalSegmentator, based on the nnU-Net framework, have shown strong performance for multi-structure CT segmentation and can be integrated into platforms such as 3D Slicer [29, 30]. However, segmentation accuracy remains critical. Errors in the delineation of the nodule or surrounding structures propagate directly to the generated model and may affect planning or guidance reliability, especially for small, low-contrast, juxta-pleural, or irregularly shaped nodules [31]. For this reason, automatic segmentation should be considered an enabling step of the workflow, but its results require manual verification and, when necessary, correction before being used for procedural support.

2.3.3 Patient-specific 3D visualization, trajectory guidance, and instrument tracking

One of the most established uses of AR and MR in image-guided interventions is the visualization of patient-specific anatomy. In this context, pre-procedural CT or Magnetic Resonance Imaging (MRI) data are processed to generate three-dimensional representations of relevant anatomical structures, target lesions, and surrounding organs. These virtual models can then be displayed through a HMD or another augmented visualization interface, allowing the operator to explore the anatomy in a more intuitive spatial format rather than mentally reconstructing it from axial, coronal, and sagittal image views [23, 21]. This can be particularly useful when the target is small, deep, poorly visible on single slices, or located near anatomical structures that should be avoided during needle insertion.

In its simplest form, patient-specific visualization acts as an interactive anatomical reference for procedural planning. Thoracic and oncological applications have used CT- or MRI-derived models to help the operator evaluate lesion location, margins, and relationships with vessels, bronchi, bones, or other relevant structures [32, 33]. In this configuration, the model supports anatomical understanding and planning, but it does not necessarily provide direct guidance for needle insertion.

A more advanced use consists in spatially aligning the patient-specific model with the physical patient or procedural scene. In this case, the virtual anatomy is not only visualized as an independent reference, but is related to the real body through a registration procedure. Registered visualization has been explored in thoracic applications using holographic or video pass-through overlays, with the aim of improving the perception of the relationship between internal anatomical structures and their external body position [27]. This step is essential when AR or

MR is intended to support procedural guidance rather than only pre-procedural understanding.

Once the model has been registered, AR and MR systems can provide different levels of trajectory support. In planned-path guidance, the system displays a target point, skin entry point, planned access route, or insertion depth, while the operator manually aligns the instrument with the virtual trajectory. In this configuration, the needle itself may not be continuously tracked: the system supports the reproduction of a pre-planned path, but the actual needle position still depends mainly on manual execution and intermittent verification. This approach can be useful when the main objective is to transfer a CT-based plan into the procedural field and reduce the cognitive gap between image-based planning and physical needle placement.

A further level of guidance is reached when the interventional instrument is also tracked. In this case, optical markers, fiducial patterns, electromagnetic sensors, or known mechanical constraints can be used to estimate the position and orientation of the needle or needle holder. The system can then display the current or projected needle path with respect to the target and surrounding anatomy, providing feedback on angular deviation, lateral displacement, insertion depth, or the relationship between the instrument and real-time imaging. Compared with planned-path guidance alone, instrument tracking can provide more dynamic feedback, although its reliability depends on tracking accuracy, marker visibility, sensor calibration, and the stability of the patient-to-image registration. The main geometric quantities involved in trajectory-based guidance are schematically represented in Figure 2.5.

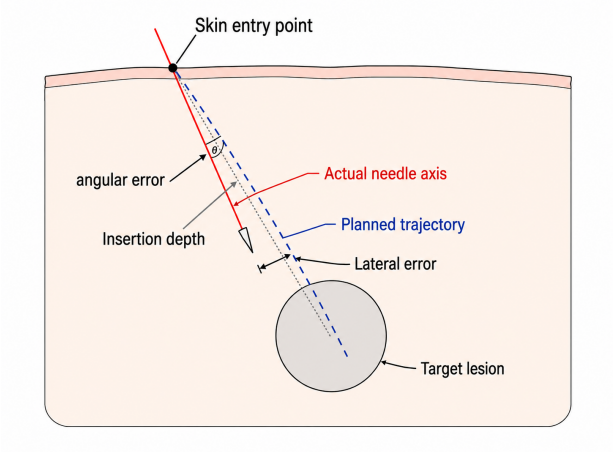

Figure 2.5: Schematic representation of trajectory guidance parameters. The planned trajectory defines the intended path from the skin entry point to the target, whereas the actual needle axis may deviate from this path during insertion. Relevant quantities include lateral error, angular error, and distance from the target.

Therefore, AR- and MR-assisted guidance systems can be seen as a continuum. At one end, they provide patient-specific 3D visualization for anatomical understanding and procedural planning. At the other end, they integrate registration, planned trajectories, real-time imaging, and, in some cases, instrument tracking to support active guidance during the procedure. Not all systems implement real-time tracking of the interventional instrument, and many rely instead on planned-path guidance without continuous feedback on needle position. Representative examples of these different approaches are summarized in Table 2.3.

Overall, these examples show that AR and MR can support image-guided interventions at different levels, from anatomical visualization to trajectory guidance and instrument tracking. Their common objective is to bring image-based information closer to the operator's field of action, reducing the need to mentally translate two-dimensional images from an external monitor into a three-dimensional manual task. However, the effectiveness of these systems depends strongly on the accuracy and robustness of patient-to-image registration, the reliability of instrument tracking, and the stability of the anatomical configuration during the procedure. These challenges are particularly relevant in thoracic interventions, where respiratory motion, lung deformation, and changes in patient position can alter the relationship between the planned model and the real anatomy. For this reason, AR- and MR-assisted systems remain promising but technically demanding solutions, whose clinical value depends on their ability to integrate visualization, registration, trajectory planning, and guidance feedback into a practical procedural workflow.

2.3.4 Challenges and limitations of mixed reality in interventional procedures

Although AR and MR systems have shown promising results for anatomical visualization, trajectory planning, and instrument guidance, their use in percutaneous interventions is still limited by several technical and clinical challenges. The final reliability of the guidance information is not determined by a single factor, but by the combination of segmentation accuracy, patient-to-image registration, anatomical motion, instrument tracking, and visualization quality.

- **Registration accuracy.** Patient-to-image registration is one of the most critical requirements for AR/MR-guided procedures. In order to use a virtual model for procedural guidance, the pre-procedural imaging dataset must be accurately aligned with the physical patient, phantom, or procedural reference. However, even small registration errors can directly affect the apparent position of the target, the planned trajectory, and the displayed relationship between the instrument and the anatomy. Therefore, registration

Table 2.3: Representative examples of patient-specific visualization, trajectory guidance, and instrument tracking approaches in AR- and MR-assisted procedures.

Study / application	Device	Registration	Tracking	Feedback provided
Perkins et al. (2020), thoracic surgical planning [32]	HoloLens	No operative patient–model registration	No real-time needle tracking	Patient-specific thoracic model, lung deflation simulation, lesion localization, and planning support
Arensmeyer et al. (2024), chest wall surgical planning [27]	Varjo XR-3 video pass-through HMD	Manual on-patient holographic alignment using anatomical landmarks and depth feedback	No real-time needle tracking	Real-time CT volume-rendered overlay, spatial assessment of internal structures relative to the body surface, and dynamic slicing/windowing
Borde et al. (2024), percutaneous needle insertion with CT-US fusion [34]	HoloLens 2 / XR90 smart goggles	Fusion between CT-derived data and ultrasound space	Electromagnetic tracking of ultrasound probe and needle	Ultrasound image, projected needle trajectory, needle tip, and CT-derived structures
Song et al. (2025), preoperative pulmonary nodule localization [35]	HoloLens	Five asymmetric labels placed on the chest wall for patient–model registration	No real-time needle tracking; label-based planned puncture guidance	Registered 3D thoracic model, target nodule, planned puncture path, and insertion-depth guidance for hookwire placement
Edgcumbe et al. (2025), DIEP flap harvest [36]	Meta Quest 3	Fiducial/tattoo-based registration with hand tracking and manual refinement	No real-time instrument tracking	Patient-specific CTA-derived vascular overlay, perforator localization, incision planning, and intraoperative marking support
Misono et al. (2026), percutaneous splenic mass biopsy [37]	HoloLens 2 / XR90 smart goggles	Registration between preprocedural CT-derived models, skin fiducials, and electromagnetic tracking space	Electromagnetic tracking of ultrasound probe and biopsy needle	Patient-specific 3D anatomy, live ultrasound visualization, tracked needle position, and combined AR/US guidance during biopsy

accuracy remains a key determinant of the reliability of AR/MR-assisted guidance.

- **Respiratory motion and deformation.** This issue is particularly relevant in thoracic interventions, where the lung is a mobile and deformable organ. The anatomical configuration during the procedure may differ from the one captured during the pre-procedural CT scan because of respiratory motion, patient positioning, gravity, tissue deformation, and needle-tissue interaction. As a result, a model that is correctly registered at the beginning of the procedure may progressively become less representative of the real anatomy during needle advancement. This limitation is especially important when guidance is based on static pre-procedural imaging, because the spatial relationship between the target lesion, pleura, vessels, bronchi, and needle may change during the intervention [21, 23].

- **Instrument tracking.** Instrument tracking is required when the interventional radiologist needs real-time feedback on the position, orientation, depth, or deviation of the needle with respect to the planned trajectory. Different solutions have been proposed, including visual patterns attached to the needle, optical markers, electromagnetic sensors, tracked ultrasound probes, or mechanically constrained needle guides. Each approach has specific advantages and limitations. Optical tracking may require a stable line of sight between the camera and the marker, electromagnetic tracking may be affected by field distortion or metallic objects, and marker-based solutions may introduce additional constraints in the sterile workflow. In addition, tracking errors, latency, loss of marker visibility, or calibration inaccuracies can reduce the reliability of the displayed trajectory during critical steps of the procedure [38, 34, 35].

- **Visualization and usability.** The main advantage of AR and MR is the possibility of bringing anatomical and trajectory information closer to the operator's field of action. However, visualization can also become a limitation if the virtual content is too dense, poorly aligned, unstable, or difficult to interpret. An excessive amount of information may obscure the patient, the skin entry region, or the instrument, increasing cognitive load instead of reducing it. Therefore, effective AR/MR guidance depends not only on registration and tracking accuracy, but also on display ergonomics, depth perception, field of view, latency, interaction modality, and the clarity of the visual feedback provided to the operator [34, 22].

Overall, these limitations show that AR and MR guidance cannot be considered only as a visualization problem. In percutaneous interventions, the final guidance

error depends on the interaction between image segmentation, registration, anatomical motion, instrument tracking, and visual presentation. This is particularly relevant for lung biopsy, where respiratory motion and tissue deformation make the translation of a static pre-procedural model into the procedural field technically demanding.

2.4 Critical summary and methodological gap

The literature reviewed in this chapter shows that several technological solutions have been proposed to support percutaneous and surgical image-guided procedures. Conventional navigation systems can improve trajectory planning, needle alignment, and instrument localization, while AR and MR approaches can bring patient-specific anatomical models and guidance information closer to the operator's field of action. However, these solutions are often focused on individual components of the workflow, such as anatomical visualization, patient-to-image registration, trajectory planning, or instrument tracking, rather than on their integration into a complete and practical procedural pipeline.

The literature includes several solutions for anatomical visualization, registration, and needle guidance, but fewer studies investigate their integration into a complete workflow running on a standalone video pass-through headset. This gap is particularly relevant for CT-guided thoracic needle insertion, where the operator must connect pre-procedural image analysis, target identification, trajectory planning, patient registration, and needle alignment within a workflow that remains compatible with the procedural environment.

This methodological gap directly motivates the design choices of the present thesis. First, many HMD-based medical AR/MR systems have relied on Microsoft HoloLens or HoloLens 2, whereas the limited future evolution of this platform motivates the investigation of alternative devices such as Meta Quest 3 [39]. Second, image processing, segmentation, planning, export, and visualization are often distributed across separate tools; therefore, this thesis develops an integrated 3D Slicer module to organize the pre-procedural planning workflow. Third, because patient-to-image registration remains a critical limitation, a QR-code-based strategy is adopted to provide a simple marker-based method for initial rigid alignment between the virtual model and the physical phantom or procedural reference. Fourth, instead of focusing only on anatomical visualization, the proposed system also includes trajectory planning and preliminary needle holder tracking, in order to explore the transition from passive visualization to active procedural guidance. Finally, while many navigation systems still require the operator to consult external monitors, the proposed MR application aims to display anatomical and trajectory information directly within the operator's field of view.

Accordingly, the goal of this thesis is not to provide clinical validation of MR-guided lung biopsy, but to demonstrate the technical feasibility of an integrated Meta Quest 3-based workflow in a controlled phantom setup. The proposed system investigates the integration of patient-specific model generation, trajectory planning, QR-code-based registration, mixed reality visualization, and preliminary instrument tracking as a basis for future developments in MR-assisted CT-guided thoracic interventions.

Chapter 3

Materials and Methods

3.1 Tools used

3.1.1 CT Dataset

For the development and testing of the proposed system, the LIDC-IDRI dataset was used [40]. This database is one of the most widely used public collections of thoracic CT scans for lung nodule analysis and computer-aided diagnosis research. It contains 1018 thoracic CT examinations acquired from different clinical institutions and annotated by four experienced thoracic radiologists. Each scan is associated with an XML file containing the final radiological annotations, including lesion location, nodule outlines when available, and subjective characteristics such as subtlety, internal structure, spiculation, lobulation, sphericity, margin, and likelihood of malignancy. In the LIDC-IDRI annotation protocol, pulmonary findings are classified into three main categories:

- Nodules with diameter ≥ 3 mm;
- Nodules with diameter < 3 mm;
- Non-nodular lesions with diameter ≥ 3 mm.

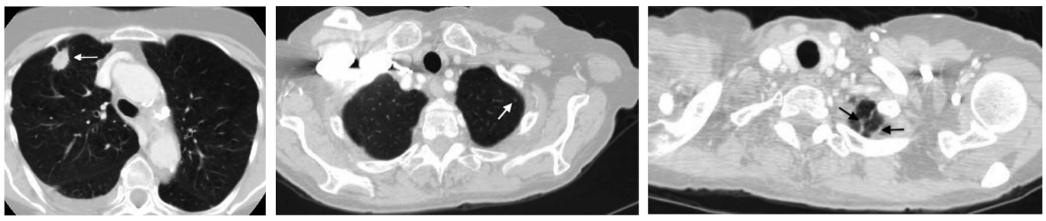

Figure 3.1: Examples of lesions satisfying the LIDC-IDRI definitions of (a) nodule ≥ 3 mm, (b) nodule < 3 mm, and (c) non-nodule ≥ 3 mm [40].

Nodules ≥ 3 mm are the most relevant category for segmentation purposes, since they are provided with radiologist-defined outlines describing their spatial extent. For this reason, the dataset represents a useful reference for testing medical image processing workflows involving lung nodule identification, segmentation, and three-dimensional model generation. The LIDC-IDRI dataset is publicly available via The Cancer Imaging Archive (TCIA) portal at the following address: `https://www.cancerimagingarchive.net/collection/lidc-idri/`.

In addition to the public LIDC-IDRI dataset, a clinical thoracic CT scan provided by interventional radiologists was also available for application-specific testing. Compared to the selected LIDC-IDRI scans used in this work, this acquisition had a slice spacing of 1 mm instead of 2.5 mm, providing a higher through-plane resolution. This additional case was useful to evaluate the workflow on data more similar to those that could be used in a real interventional planning scenario. Moreover, the higher resolution can support the generation of more detailed 3D models, which is particularly beneficial for the visualization of small anatomical structures.

3.1.2 3D Slicer

3D Slicer is a free, open-source software platform for the visualization, processing, segmentation, registration, and analysis of medical images and three-dimensional data [41]. Its modular architecture makes it particularly suitable for research workflows, since its functionalities can be extended through native modules, installable extensions, and custom modules developed in Python or C++ [42]. For this reason, 3D Slicer is widely used not only for medical image analysis, but also for surgical planning and image-guided procedure research.

In this thesis, 3D Slicer was used as the platform on which a custom module, named *Lung Nodule Biopsy Planner*, was developed to collect within a single interface the main operations required in the first part of the mixed-reality workflow, from the import of the thoracic CT scan to the generation of the anatomical models and planned biopsy trajectory to be used in the Meta Quest 3 application.

Several existing 3D Slicer modules and extensions were integrated into this custom workflow. TotalSegmentator was used for automatic segmentation of thoracic anatomical structures [29] and is described separately in the following subsection. The Percutaneous Approach Analysis module was used to support trajectory planning by computing candidate needle paths from a selected target point, skin surface model, and obstacle model [43]. In addition, the PortPlacement module was used for visual verification of the planned instrument orientation, while OpenIGTLink was used to enable real-time communication between 3D Slicer and the mixed-reality application [44], with the goal of visualizing the slices corresponding to the section cut by a moving plane sliding across the model. The

detailed implementation of the custom module, including the optimization of the trajectory analysis step, is described in the following sections.

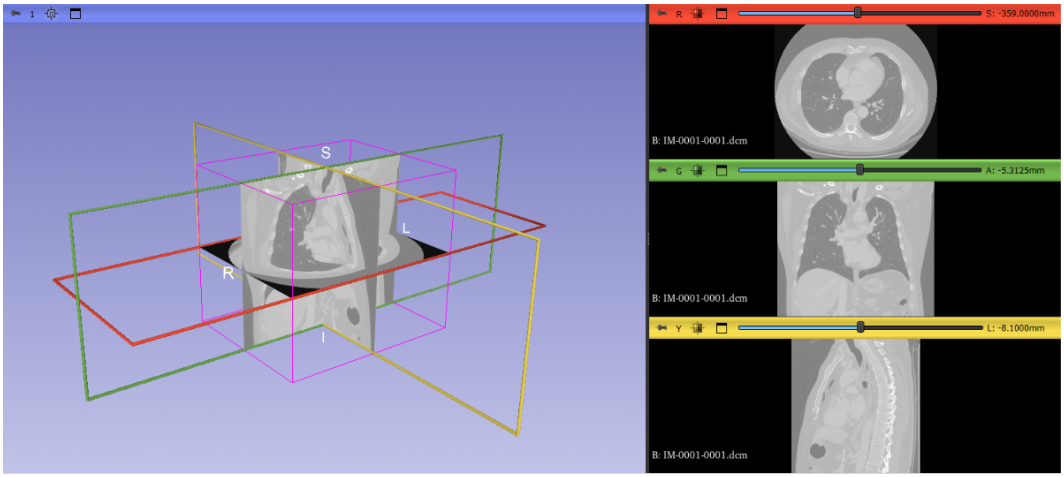

Figure 3.2: Standard 3D Slicer views for thoracic CT volume analysis: 3D view with slicing planes (left) and 2D axial, coronal, and sagittal views (right).

3.1.3 TotalSegmentator

TotalSegmentator was used as the automatic segmentation tool within the 3D Slicer workflow. It is an open-source deep learning-based framework for anatomical segmentation in CT images, based on the nnU-Net architecture. In the original validation study, the model was trained on a heterogeneous dataset of 1204 CT examinations and provided segmentations of 104 anatomical structures, reporting a mean Dice similarity coefficient of 0.943 for the high-resolution model [29]. In the current public implementation, the default `total` task includes 117 main CT classes, while additional task-specific models are available for selected anatomical regions [45].

The choice of TotalSegmentator was motivated by its broad anatomical coverage, public availability, integration within 3D Slicer, and the availability of task-specific models relevant to thoracic biopsy planning. In this thesis, four tasks were selected: `total`, `lung_nodules`, `lung_vessels`, and `body`. These tasks allowed the extraction of the main thoracic anatomical structures, lungs and candidate pulmonary nodules, pulmonary vessels and airways, and the external body surface.

Since no manually annotated ground truth was available in this work, the segmentations generated by TotalSegmentator were not quantitatively validated. They were instead visually inspected and used as intermediate anatomical models for planning, visualization, and mixed-reality scene preparation.

3.1.4 Unity and Meta XR SDK

Unity was used as the main development environment for the mixed-reality application running on Meta Quest 3. Unity is a real-time 3D development platform that allows the creation of interactive scenes, the management of three-dimensional assets, and the implementation of application logic through C# scripts [46]. In this thesis, Unity version 6000.2.12f1 was used to develop the application responsible for importing, displaying, manipulating, and registering the patient-specific thoracic model generated in 3D Slicer.

The Unity project was configured for Meta Quest development using the Meta XR SDK All-in-One package, version 201.0.0. This package includes the core components required to build applications for Meta XR devices, including headset tracking, rendering support, input management, passthrough visualization, and mixed-reality interaction features [47, 48]. In this context, Meta Quest 3 was used as a standalone video pass-through headset, allowing virtual anatomical content to be displayed and interacted with directly in the real environment rather than on an external monitor.

Several custom C# scripts were developed to implement the specific functionalities required by the proposed workflow. These included the management of the imported thoracic model, user interaction with the virtual anatomy, visibility and manipulation controls, communication with the local server used to retrieve the exported GLB file, and the logic required for registration and guidance. The glTFast package was integrated to support runtime loading of the GLB model exported from 3D Slicer, allowing the application to automatically import the most recently generated patient-specific scene through a local web server [49].

Meta Mixed Reality Utility Kit (MRUK) was also integrated into the project to support mixed-reality-specific functionalities. In particular, MRUK was used for QR code detection during the registration phase. In the Meta framework, detected QR codes can be treated as trackable objects, providing their spatial pose and encoded payload within the Unity scene [50]. This allowed the application to use physical QR markers as references for aligning the virtual thoracic model with the real environment.

During development, the Meta XR Simulator was used to test and debug the application directly from the Unity Editor before deploying it to the Meta Quest 3. This simulation environment (Figure 3.3) made it possible to verify basic interactions, scene behavior, and mixed-reality logic on the workstation, reducing the need for repeated builds and allowing faster debugging during the early development phases [51]. Final testing and validation were then performed on the physical headset.

Overall, Unity represented the second main component of the proposed workflow. While 3D Slicer was used for medical image processing, segmentation, preliminary

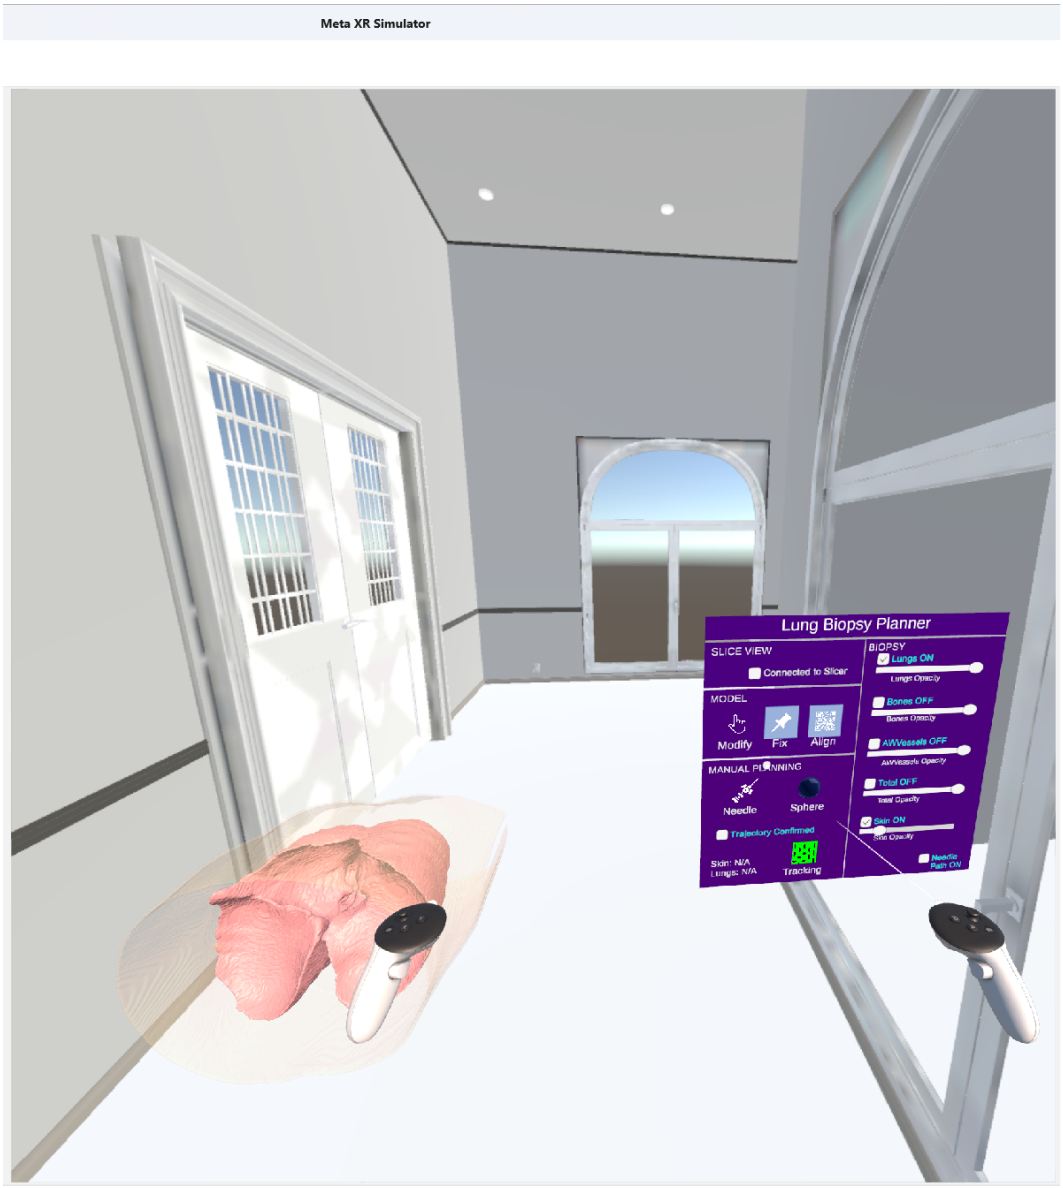

Figure 3.3: Meta XR Simulator running in Unity Editor, simulating the Meta Quest 3 environment.

trajectory planning, and model preparation, Unity was used to transform these outputs into an interactive mixed-reality application running on Meta Quest 3.

3.1.5 Meta Quest 3

Meta Quest 3 was selected as the head-mounted display for the mixed-reality application developed in this thesis. It is a standalone virtual- and mixed-reality headset, capable of running applications directly on the device, but it can also be connected to an external workstation through Meta Horizon Link [52, 53]. In this work, the headset was mainly used through a wired USB-C connection during development and visualization tests. This configuration allowed the Unity application to be executed on the workstation while being visualized and interacted with through the headset, taking advantage of the PC GPU for rendering. In the tests performed in this work, this configuration allowed more detailed anatomical models to be displayed with less aggressive mesh decimation and with a larger number of visible anatomical structures. Native execution on the headset was also considered for testing device-side mixed-reality functionalities, as described in the implementation sections.

Figure 3.4: Meta Quest 3 headset.

Meta Quest 3 is based on video pass-through technology. In this architecture, the real environment is captured by external cameras and displayed to the user through the headset displays, while virtual content is rendered and blended into the same visual scene. This approach differs from optical see-through headsets, where the real world is observed directly through transparent optics. In the context of this thesis, video pass-through was used to overlay the thoracic model, the planned trajectory, and registration references within the user's field of view, enabling the anatomical models generated in 3D Slicer to be visualized in relation to the surrounding physical environment.

The device integrates a Qualcomm Snapdragon XR2 Gen 2 platform, pancake lenses, high-resolution displays, inside-out tracking, hand tracking, color pass-through cameras, and depth sensing for room mapping and mixed-reality applications [52]. These features make it suitable for interactive applications requiring spatial awareness, real-time rendering, and user interaction with three-dimensional virtual content. The main technical characteristics of the headset relevant to the proposed workflow are summarized in Table 3.1.

From a sensing point of view, the components most relevant to this thesis are the RGB cameras used for color video passthrough, the depth sensing system used for spatial understanding, and the inside-out tracking system used to estimate the pose of the headset and controllers relative to the environment. These sensors provide the basis for displaying virtual anatomical content in a stable position within the physical scene and for implementing marker-based registration strategies.

For the application developed in this work, the most relevant features were the video pass-through system, the inside-out tracking capabilities, the support for mixed-reality development through the Meta XR SDK, and the possibility of testing the application both on the headset and through a PC-based development workflow. The pass-through system enabled virtual anatomical content to be displayed in relation to the physical environment, while inside-out tracking provided the spatial reference required to maintain virtual objects in stable positions relative to the user and the surrounding scene. The Meta XR SDK and MRUK were then used within Unity to implement interaction and registration functionalities, including QR-based alignment.

In this thesis, Meta Quest 3 was therefore used as the final visualization and interaction platform of the workflow. The pre-procedural CT data were processed in 3D Slicer, exported as a three-dimensional scene, loaded into Unity, and displayed on the headset as an interactive mixed-reality model.

3.1.6 OpenIGTLink

OpenIGTLink is an open network communication protocol designed for image-guided therapy environments. It was introduced to support standardized communication between software applications, imaging systems, tracking devices, and other hardware components used in image-guided interventions [56]. The protocol is particularly suitable for research systems that require interoperability between heterogeneous platforms, since it supports the exchange of different types of data, including images, spatial transformations, tracking information, and device status messages.

In 3D Slicer, OpenIGTLink communication is handled by the OpenIGTLinkIF module, which enables network communication with external applications and

Table 3.1: Main technical specifications of Meta Quest 3 relevant to the proposed mixed-reality workflow [52, 53, 54, 55].

Category / feature	**Specification**
Device type	Standalone virtual and mixed-reality head-mounted display
Platform / OS	Meta Quest platform; Android-based operating system
Processor (SoC)	Qualcomm Snapdragon XR2 Gen 2
System memory	8 GB RAM
Storage configuration used	512 GB flash storage
Display type	Binocular LCD
Display resolution	2064 × 2208 pixels per eye
Refresh rate	Up to 120 Hz
Optics	Pancake lenses
IPD adjustment range	58–71 mm (mechanical adjustment)
Field of view	Approx. 110° horizontal, 96° vertical
Mixed-reality sensing	2 RGB color cameras and IR depth projector for color video passthrough and scene understanding
Passthrough Camera API stream	Camera frames accessible to applications at 60 Hz, with an image capture latency of approximately 20–40 ms and a maximum resolution of 1280 × 1280 pixels
Tracking modalities	Inside-out 6DoF headset tracking, Touch Plus controller tracking, and hand tracking
Controllers	Meta Quest Touch Plus controllers (6DoF)
Audio	Integrated stereo spatial audio; built-in microphone
Connectivity	Wi-Fi 6E, Bluetooth; USB-C port for power, data, and Meta Horizon Link; Air Link for wireless PC VR streaming
Device form factor	Approx. 512 g; plastic/foam construction; fabric head strap
Battery life	Approximately 2.2 hours of active use, depending on application and usage conditions
Charging time	Approximately 2.3 hours (from empty to full charge)

devices [57]. It supports the import and export of image data and linear transformations, making it suitable for synchronizing Slicer with external visualization or navigation systems.

In this thesis, OpenIGTLink was used as the communication layer between 3D Slicer and the Unity application. The objective was not to transfer the complete three-dimensional anatomical model, which was loaded separately in Unity as a GLB file, but to exchange dynamic information related to slice visualization. In particular, the application transmitted the transformation associated with a mobile cutting plane, while 3D Slicer generated the corresponding axial CT reslice and sent it back to Unity for visualization in the mixed-reality scene.

This approach was inspired by previous work demonstrating real-time integration between 3D Slicer and a Unity-based augmented-reality application through OpenIGTLink, where geometrical transformations and CT image messages were exchanged between the two platforms to display resliced CT images within a head-mounted display [58]. In the present work, a similar communication principle was adapted to the Meta Quest 3 application and to the thoracic biopsy planning scenario.

The Volume Reslice Driver module was used to support the reslicing operation inside 3D Slicer. This module allows a slice plane to be driven by a linear transformation, enabling the generation of a sectional image from a volumetric dataset according to the position and orientation of a selected reference transform [59]. In this thesis, this functionality was used to obtain the CT slice corresponding to the position of the moving plane defined in the mixed-reality environment. The detailed implementation of the OpenIGTLink connection, including message handling and synchronization between Unity and 3D Slicer, is described in the following sections.

3.1.7 Cylindrical marker for instrument tracking

To enable preliminary instrument tracking, a cylindrical visual marker inspired by the design proposed by Cartucho et al. was considered [60]. The marker is conceived as a fiducial pattern for tracking elongated cylindrical instruments, such as needles, cannulas, or trocars, by providing a visual reference from which the spatial pose of the instrument can be inferred.

The intended configuration of the system involves the design and 3D printing of a dedicated support surrounding the cannula or biopsy needle. This support would be positioned on the portion of the instrument that remains outside the patient during the procedure, without interfering with the part inserted into the body. In this way, the marker could be rigidly coupled to the instrument while preserving a known geometric relationship with its axis. Once this relationship is defined, the pose of the marker can be used as an indirect estimate of the pose of the instrument itself.

The marker pattern consists of a green cylindrical surface containing black elliptical features arranged according to a predefined binary code. Each ellipse encodes a binary value according to its orientation, so that differently oriented ellipses represent the two possible states of the code. This type of pattern is particularly suitable for cylindrical objects, since it can be distributed around the instrument and observed from different viewpoints. The green background provides visual contrast with the black features, making the marker easily recognizable with respect to the surrounding scene.

In the preliminary phase of this work, before designing and fabricating the dedicated 3D-printed support, the marker was applied to a pen with a diameter of 13.38 mm used as a cylindrical surrogate of the biopsy instrument. This simplified prototype allowed the visual appearance of the marker, its arrangement around a cylindrical object, and its potential integration into the mixed-reality workflow to be assessed without requiring a clinical instrument. A photograph of this preliminary marker prototype is shown in Figure 3.5.

Figure 3.5: Preliminary prototype of the cylindrical visual marker inspired by Cartucho et al. [60]. In this early implementation, the marker pattern was applied to a pen used as a cylindrical surrogate of the biopsy instrument.

In the context of this thesis, the marker was therefore used as the physical tracking target for a prototype vision-based instrument tracking strategy. Its role was to provide a recognizable and geometrically defined reference on the instrument, enabling the mixed-reality application to estimate the orientation of the tracked object and compare it with the planned trajectory. This information is relevant for the visualization of guidance-related quantities, such as the angular deviation between the instrument and the planned path and the lateral deviation from the desired trajectory.

It should be noted that, in this work, the marker-based tracking strategy was

considered as a prototype component of the system. Its purpose was to support the technical integration of instrument tracking within the mixed-reality workflow, rather than to provide a clinically validated navigation solution.

3.1.8 Flask server and OpenCV

A Flask server was used as an auxiliary Python-based component to support the communication between the Meta Quest 3 application and the computer vision tools adopted for marker-based instrument tracking. Flask is a lightweight web framework for Python that allows the creation of HTTP endpoints through which external applications can exchange data with a Python process [61]. In this thesis, it provided a simple communication interface between the Unity application and the Python environment used for processing the visual marker described in the previous subsection.

OpenCV was selected as the computer vision library for the marker tracking component. It is an open-source library that provides tools for image processing, camera calibration, feature analysis, and geometric pose estimation [62]. Within the proposed system, OpenCV represented the software basis for interpreting the cylindrical marker and estimating the pose information required by the mixed-reality application.

The combination of Flask and OpenCV made it possible to keep the mixed-reality visualization logic in Unity while relying on Python for the computer vision component. The detailed implementation of image acquisition, marker processing, pose estimation, and coordinate conversion is described in the following sections.

3.2 Pre-operative planning workflow in 3D Slicer

The first part of the proposed workflow was performed in 3D Slicer through a custom module developed for this thesis, named *Lung Nodule Biopsy Planner*. The module was designed to centralize the main pre-operative planning operations required for CT-guided lung biopsy support within a single interface. These operations include loading and inspecting the thoracic CT volume, generating patient-specific anatomical segmentations, defining the target nodule and the anatomical structures to be avoided, computing candidate percutaneous trajectories, visually verifying the selected path, and preparing the resulting scene for mixed-reality visualization on Meta Quest 3.

The rationale behind this custom module was to reduce workflow fragmentation. Instead of requiring the operator to move manually between several independent 3D Slicer modules, the most relevant steps were integrated into a sequential interface tailored to the biopsy planning scenario. The module incorporated

existing tools and extensions, including TotalSegmentator for automatic anatomical segmentation, the Markups system for target definition, Percutaneous Approach Analysis for trajectory computation, PortPlacement for visual verification of the planned instrument orientation, and OpenIGTLink for communication with the Unity application. Figure 3.6 summarizes the complete pre-operative planning workflow implemented in the custom 3D Slicer module.

3.2.1 CT loading and visualization

The input of the planning workflow was a thoracic CT volume in NIfTI format. After loading the volume into 3D Slicer, the scan could be inspected in the conventional axial, coronal, and sagittal views, together with the three-dimensional view. Window width and window level controls were included in the custom interface to adjust image contrast and brightness according to the anatomical structures of interest, particularly the lung parenchyma, pulmonary vessels, ribs, and target nodule.

This preliminary visualization step was necessary to verify that the input volume was correctly loaded, to inspect the anatomical region of interest, and to support the subsequent manual selection of the target point. No image preprocessing aimed at modifying the original diagnostic information was performed at this stage; the volume was used as the reference dataset for segmentation, trajectory planning, and slice visualization in the mixed-reality environment.

3.2.2 Automatic segmentation and anatomical model generation

After loading the thoracic CT volume, automatic segmentation was performed using the TotalSegmentator tasks selected for the workflow. The `total` task was used to obtain the main anatomical structures of the thoracic region, with particular interest in the rib cage and other structures useful for spatial orientation. The `lung_nodules` task was used to extract the lungs and candidate pulmonary nodules. The `lung_vessels` task was used to segment pulmonary vessels and airways, while the `body` task was used to generate the external body surface.

The resulting segmentations were visualized in the 3D Slicer scene (Figure 3.8) and inspected before being used in the following planning steps. The skin surface obtained from the `body` task was used as the candidate entry region for the percutaneous needle. Selected structures, such as ribs, vessels, and airways, were combined to define the obstacle model used during trajectory computation. The segmented lungs, nodules, vessels, airways, bones, and skin surface were also included in the patient-specific three-dimensional scene prepared for export to the Meta Quest 3 application.

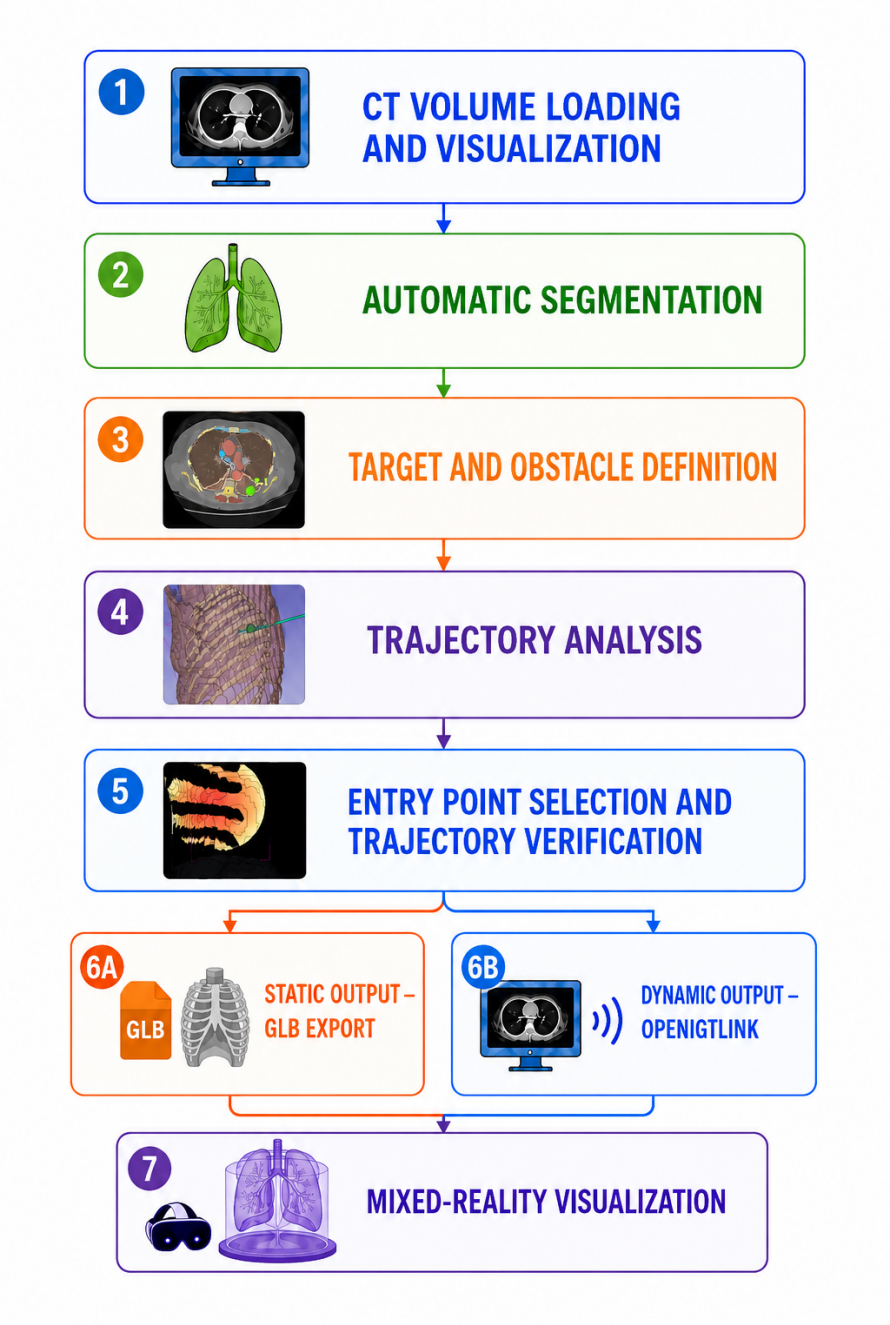

Figure 3.6: Pre-operative planning workflow implemented in the custom 3D Slicer module.

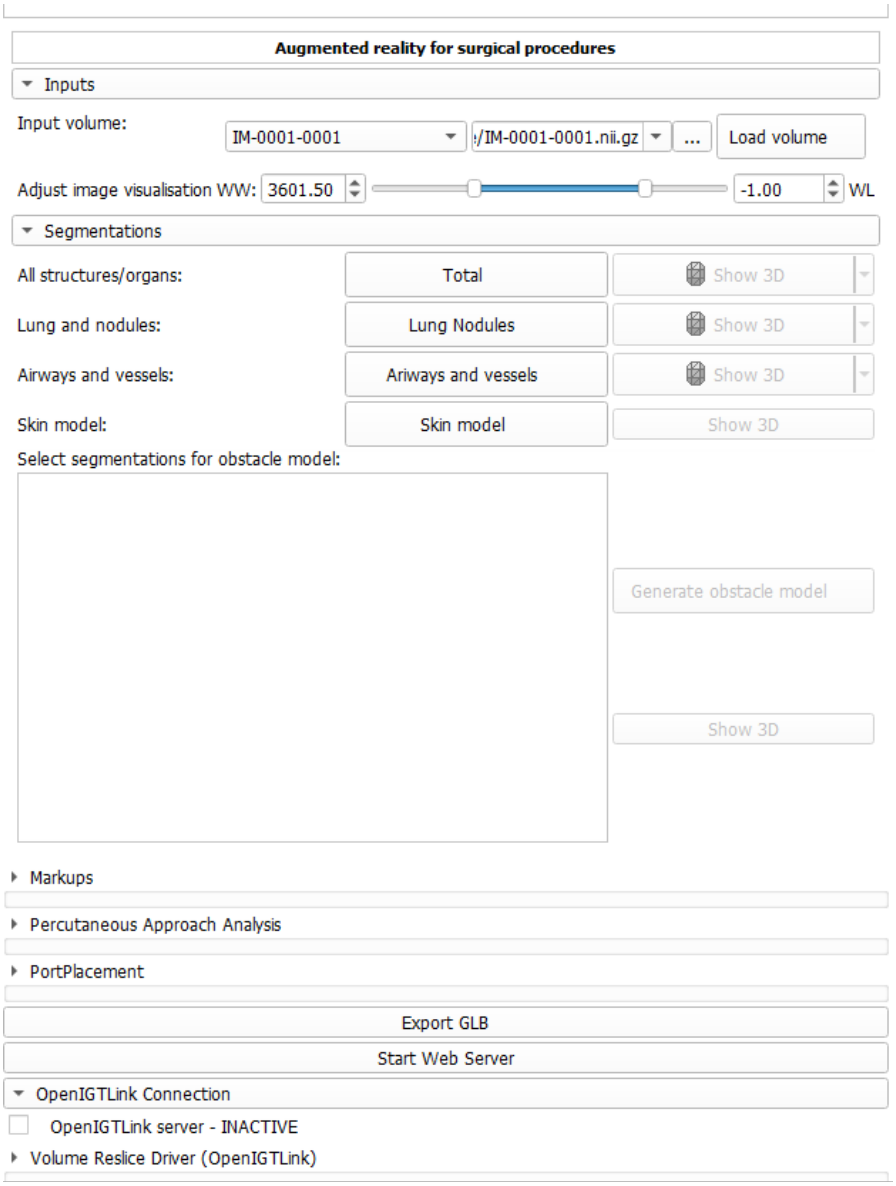

Figure 3.7: Custom 3D Slicer module interface for pre-operative planning. The interface centralizes the main operations required for thoracic biopsy planning, including CT loading, segmentation, target definition, trajectory analysis, and scene export for mixed-reality visualization.

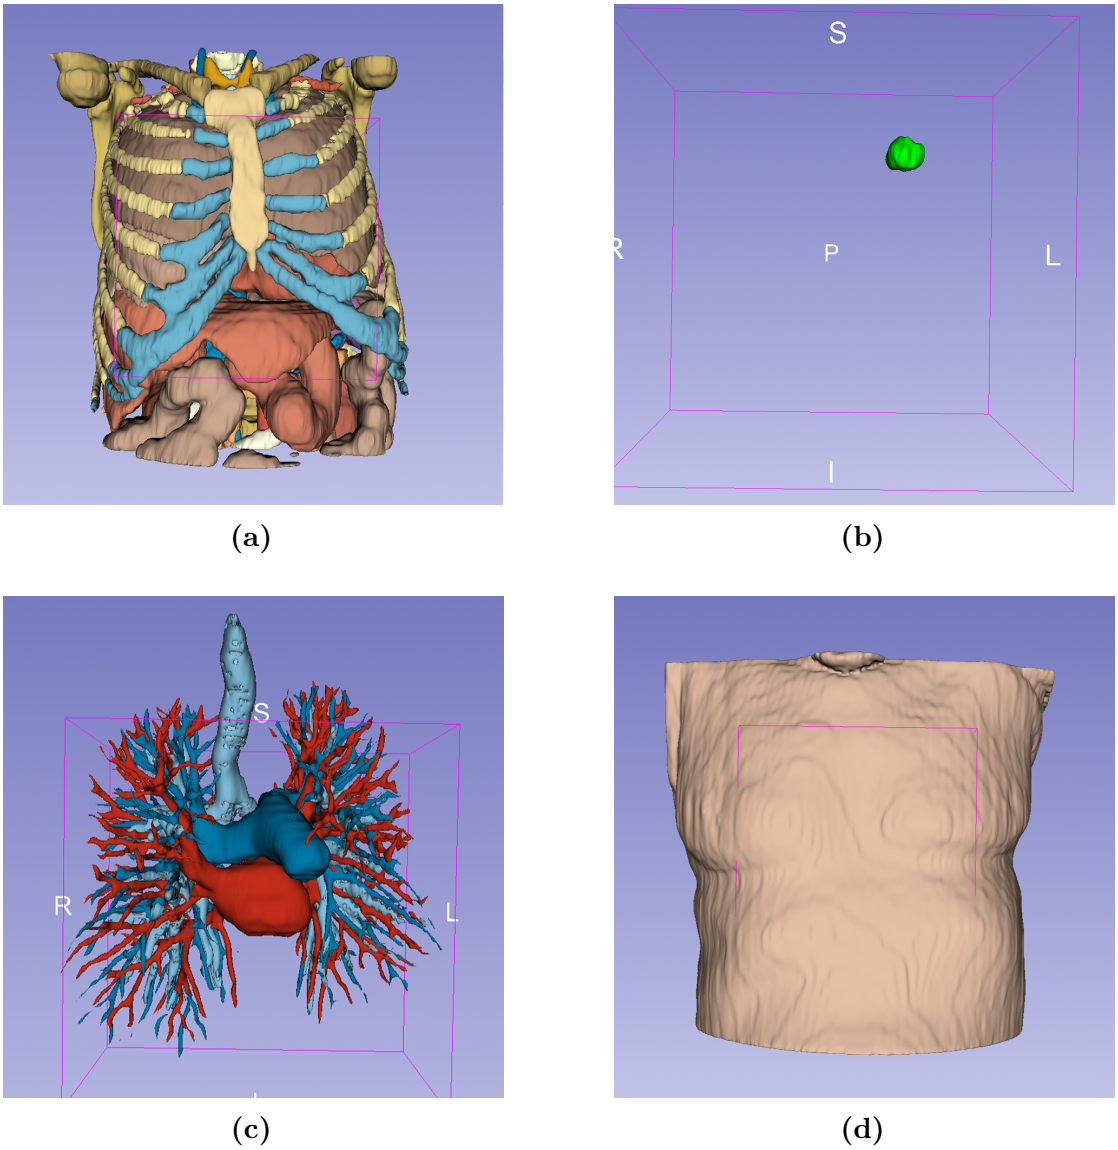

(a) **(b)**

(c) **(d)**

Figure 3.8: Visualization in 3D Slicer of the segmentation outputs generated with the four TotalSegmentator tasks used in the planning workflow: (a) main thoracic anatomical structures obtained with the `total` task; (b) candidate pulmonary nodule obtained with the `lung_nodules` task; (c) pulmonary vessels and airways obtained with the `lung_vessels` task; and (d) external body surface obtained with the `body` task.

3.2.3 Target definition, obstacle model, and trajectory analysis

After the generation of the anatomical segmentations, the target nodule was defined within 3D Slicer using the Markups system, as shown in Figure 3.9. The target was selected manually on the CT slices by placing a fiducial point at the desired biopsy target location. This point represented the final endpoint of the planned needle trajectory and was used as the main target input for the trajectory analysis step.

The anatomical structures to be avoided during needle insertion were then selected from the available segmentations and combined into an obstacle model. In the proposed workflow, this model could include structures such as ribs, pulmonary vessels, and airways, depending on the anatomical configuration of the case and on the structures considered relevant for safe trajectory planning. The external body surface obtained from the skin segmentation was used as the candidate entry region for the percutaneous approach. An example of the resulting obstacle model is shown in Figure 3.10.

Trajectory analysis was performed using the Percutaneous Approach Analysis module integrated into the custom planning interface. The module was used to configure the geometric inputs required for the computation and to visualize the resulting candidate trajectories. As shown in Figure 3.11, the interface was divided into a parameter section, used to define the input models, and an outcome section, used to inspect the results of the analysis.

In the parameter section, the user selected the target fiducial point (Figure 3.9), the output fiducial list (the skin entry point in our case), the skin surface model, and the obstacle model. The target point represented the final endpoint of the planned needle trajectory, while the skin model defined the possible entry surface. The obstacle model contained the anatomical structures to be avoided during needle insertion. Opacity sliders were provided for both the skin and obstacle models, allowing the user to adjust their visualization during the planning phase. Once these inputs were selected, the analysis was started directly from the same interface.

Based on these inputs, the module computed candidate linear paths connecting the skin surface to the target point while checking whether each path intersected the selected obstacle model. The outcome section allowed the user to visualize the color-mapped skin surface, the complete set of feasible candidate paths, and specific trajectories of interest. In particular, the interface displayed the accessibility score, the total number of generated paths, the longest path, the shortest path, and an individual candidate path that could be selected for inspection. The corresponding path length was reported in millimetres. In this implementation, the shortest path was generally considered the most suitable candidate for biopsy.

The original Percutaneous Approach Analysis workflow was based on repeated line-obstacle intersection tests. For each candidate entry point on the skin surface,

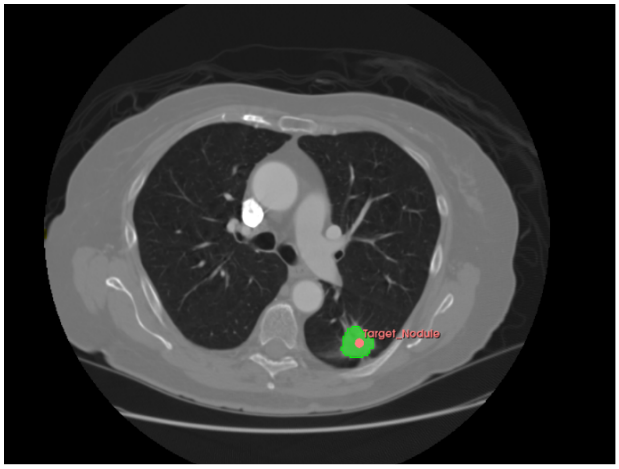

Figure 3.9: Manual definition of the target nodule in 3D Slicer using the Markups system. A fiducial point is placed at the desired biopsy target location on the CT slices.

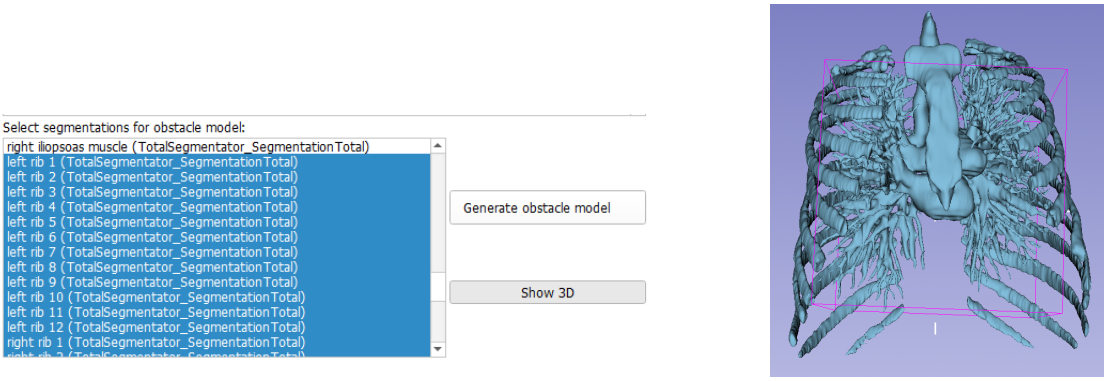

Figure 3.10: List of the anatomical structures available for selection (left) and an example of the resulting obstacle model (right) generated by combining the segmented ribs, pulmonary vessels, and airways. This model was used during trajectory analysis to identify feasible paths from the skin surface to the target nodule.

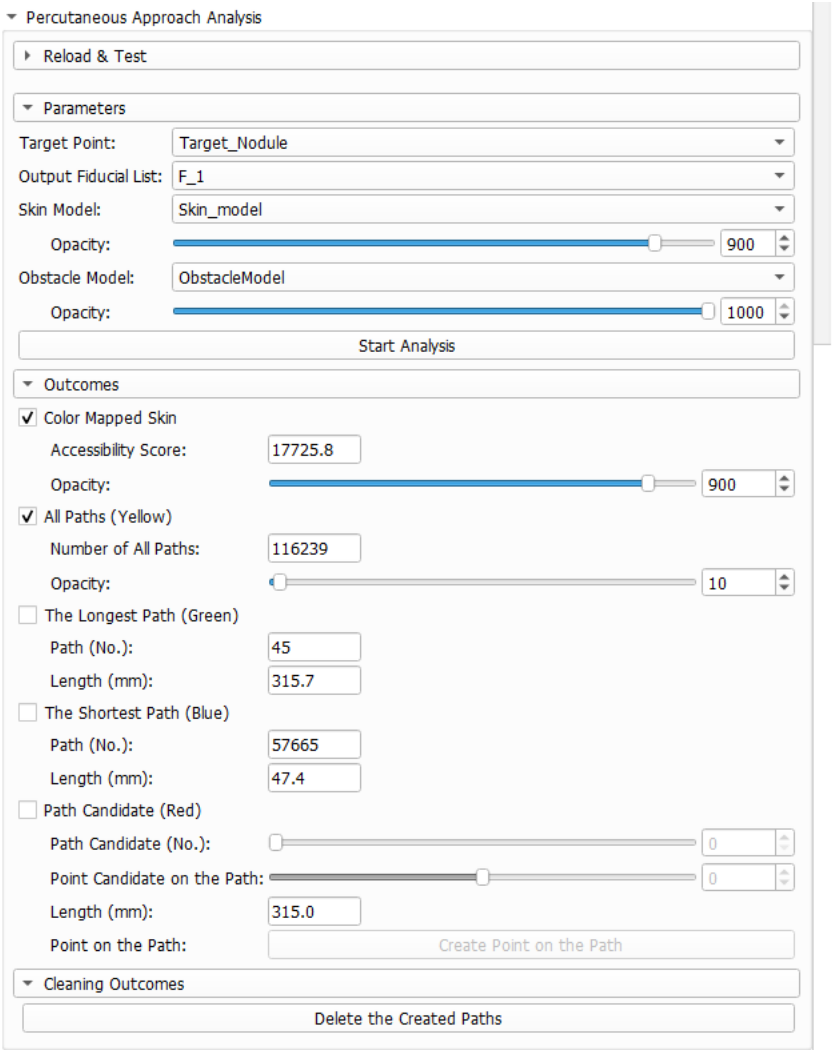

Figure 3.11: Percutaneous Approach Analysis interface integrated into the custom planning workflow. The parameter section allows the user to select the target point, skin surface, and obstacle model, while the outcome section displays the accessibility map, candidate paths, shortest and longest trajectories, and selected trajectory information.

a line segment connecting the skin point to the target was tested against the obstacle model. If the segment intersected the obstacle mesh, the corresponding trajectory was discarded; otherwise, it was stored as a feasible candidate path. A similar approach was used for the accessibility map, where each triangle of the skin mesh was evaluated by testing the line connecting its centroid to the target point. Although this approach already relied on a spatial locator, the repeated evaluation of individual points and triangles resulted in long computation times for detailed anatomical models.

To make the trajectory analysis more suitable for an interactive planning workflow, the original implementation was modified by integrating Open3D into the 3D Slicer module to speed up these computations [63]. In the optimized implementation, the obstacle model was converted into a triangular mesh and inserted into an Open3D raycasting scene. Candidate trajectories from the skin surface to the target were then represented as rays and evaluated in batch, rather than being tested one at a time.

The same strategy was applied to the accessibility mapping. The skin surface was triangulated, and for each triangle the centroid and area were computed. Rays from the triangle centroids to the target were then tested against the obstacle model. A skin region was considered accessible when the corresponding ray reached the target without intersecting any obstacle and when the path length was below the predefined maximum distance of 130 mm. The accessibility score was computed as a distance-weighted sum of the accessible triangle areas, while the corresponding scalar values were assigned to the skin mesh to generate the colorimetric accessibility map shown in Figure 3.12. In this map, more saturated colors toward red indicate shorter accessible paths, whereas progressively lighter colors correspond to longer ones.

This modification did not change the conceptual output of the Percutaneous Approach Analysis module. The module still produced candidate paths, shortest and longest trajectories, an accessibility score, and a color-mapped skin surface. However, the computational strategy was changed from repeated per-point and per-triangle intersection tests to a more efficient raycasting-based implementation. This optimization was introduced because the original workflow required approximately 12 minutes for the considered planning case, limiting its practical use during iterative trajectory exploration.

After selecting the planned entry point, the trajectory was visually verified using the PortPlacement module integrated into the *Lung Nodule Biopsy Planner* interface. The target nodule and the selected skin entry point were used to define the planned needle direction, allowing the virtual instrument to be oriented along the selected path. This step provided a visual check of the planned trajectory with respect to the segmented anatomy before preparing the scene for mixed-reality

visualization, as shown in Figure 3.13.

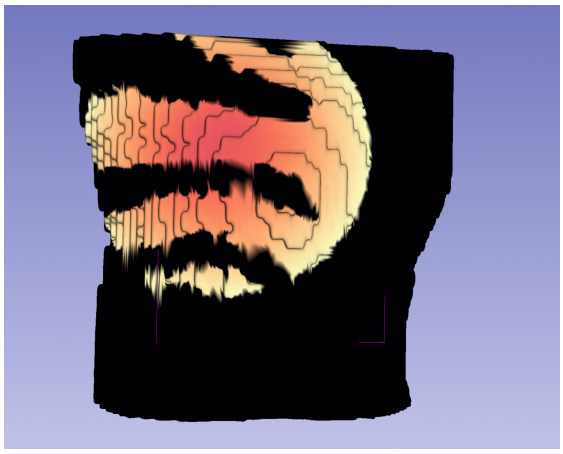

Figure 3.12: Colorimetric accessibility map generated on the skin surface. Redder regions indicate shorter accessible paths to the target, whereas lighter regions correspond to longer paths.

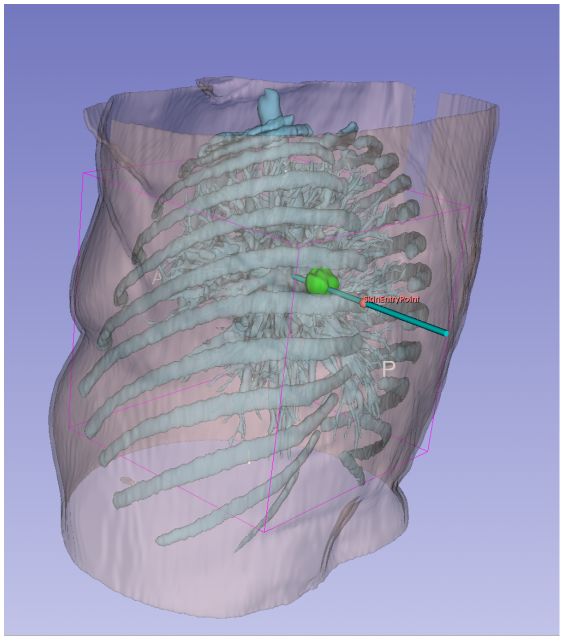

Figure 3.13: Visual verification of the planned trajectory in 3D Slicer. The virtual instrument is oriented along the selected path from the skin entry point to the target nodule, allowing a visual check of the planned trajectory.

3.2.4 OpenIGTLink configuration for slice synchronization

Before exporting the planned scene, the OpenIGTLink connection was configured to support the synchronization of CT slice visualization between 3D Slicer and the Unity application. While the patient-specific anatomical model was exported as a static GLB file, the CT slice visualization required a dynamic communication channel, since the slice displayed in the mixed-reality scene had to correspond to the position of a movable cutting plane controlled in Unity.

For this purpose, the OpenIGTLink server was activated directly from the *Lung Nodule Biopsy Planner* interface, and the Volume Reslice Driver module was used to link the reslicing operation in 3D Slicer to the transform associated with the moving plane. In this configuration, Unity transmitted the position of the virtual slicing plane to 3D Slicer, while 3D Slicer generated the corresponding CT reslice and sent it back to Unity through OpenIGTLink [57, 59].

When the OpenIGTLink connection was enabled, the input CT volume was automatically recentered with respect to the origin of the 3D Slicer coordinate system. This step was introduced to maintain a consistent spatial correspondence between the CT volume, the reslicing plane, and the coordinate frame used by the Unity application. By aligning the geometric center of the volume with the origin, the position of the moving plane in Unity could be associated more directly with the corresponding slice position in 3D Slicer.

3.2.5 Export for mixed-reality visualization

After trajectory verification and slice synchronization setup, the patient-specific thoracic scene was exported in GLB format. The export procedure was designed to include the anatomical structures and planning elements required for visualization in Unity, while excluding auxiliary elements used only during the planning phase. In particular, the obstacle model used for trajectory computation was not exported as a separate visualization object, since its role was limited to the analysis of candidate paths. The exported scene included the selected patient-specific anatomical models, the skin surface, the target-related structures, and the planned instrument representation generated during the verification step.

To reduce the computational load of the mixed-reality application, mesh decimation was applied before export. This step reduced the number of triangles of the generated surface models while preserving their overall anatomical shape. Different reduction levels were used according to the type of structure: larger and smoother structures, such as the lungs and skin surface, could be simplified more aggressively, while more complex structures, such as airways, vessels, ribs, and the planned instrument, required a more conservative reduction to preserve their visual and spatial information.

Two GLB versions of the same planned scene were generated. The first one was a full-resolution version, intended for use in the Unity Editor or through Meta Horizon Link, where rendering could rely on the workstation GPU. The second one was a lighter version, specifically prepared for standalone execution on the Meta Quest 3. In this version, only the anatomical structures required for the mixed-reality workflow were retained, and mesh decimation was applied more aggressively to reduce the computational load on the headset.

During export, the selected models and segmentation-derived surfaces were converted into triangular meshes, cleaned to remove invalid geometry, grouped into anatomical categories, and saved as GLB files. The use of the GLB format allowed the three-dimensional scene to be stored in a compact binary file, including geometry and color information. After export, a local web server was started from within the 3D Slicer module to make the generated files available to the Unity application. This avoided the need for manual file transfer and allowed the mixed-reality application to load the most recent planning scene directly from the workstation.

The exported GLB files, together with the OpenIGTLink-based slice synchronization setup, represented the final output of the pre-operative planning workflow in 3D Slicer. These outputs provided the anatomical, planning, and slice-visualization information required by the following mixed-reality phase, where the patient-specific model was imported, displayed, manipulated, and registered within the Meta Quest 3 environment.

3.3 Mixed-reality application in Unity for Meta Quest 3

The second part of the proposed workflow was implemented in Unity as a mixed-reality application for Meta Quest 3. While 3D Slicer was used for medical image processing, segmentation, trajectory planning, and scene preparation, the Unity application was responsible for importing the patient-specific model, displaying it in the mixed-reality environment, managing user interaction, registering the virtual anatomy to the physical setup, and providing preliminary guidance and tracking functionalities.

The application received two main types of information from the 3D Slicer workflow. The first was the static patient-specific anatomical scene exported as a GLB file, including the segmented anatomical structures and the planned trajectory. The second was dynamic slice information transmitted through OpenIGTLink, which allowed the visualization of CT slices corresponding to a movable cutting plane controlled in the mixed-reality environment.

The Unity application was therefore designed as the interactive component of the system, integrating anatomical visualization, model interaction, automatic QR-based registration, dynamic CT slice visualization, virtual trajectory guidance, and preliminary tracking of the real needle.

3.3.1 Automatic import of the patient-specific 3D model

The first step in developing the mixed-reality application in Unity was the import of the patient-specific thoracic scene generated in 3D Slicer. The scene was exported from the *Lung Nodule Biopsy Planner* module in GLB format and made available through the local web server started by 3D Slicer. This setup enabled the application to automatically retrieve and load the model at startup, eliminating the need for manual file transfer between the two environments.

The GLB file was loaded at runtime using the glTFast package, ensuring that Unity always accessed the most recent version of the anatomical scene. Once imported, however, the model could not be used directly, as 3D Slicer and Unity rely on different spatial conventions.

In 3D Slicer, geometry is defined in the RAS patient coordinate system, where the axes correspond to Right, Anterior, and Superior directions, and all distances are expressed in millimetres. Unity, on the other hand, uses a left-handed coordinate system with X pointing to the right, Y upward, and Z forward, and interprets one unit as one metre. Because of these differences, both the scale and the orientation of the model needed to be adjusted. In particular, to account for the change in handedness between the two systems, a flip of the X and Z axis was applied, ensuring that the imported geometry was not mirrored and remained correctly oriented in the Unity scene.

First, the model was uniformly scaled by a factor of 0.001 to convert measurements from millimetres to metres. Then, a coordinate transformation was applied to map the RAS axes onto Unity's reference frame and to account for the change in handedness. This step ensured that the anatomical structures appeared with the correct orientation and were not mirrored or misaligned in the mixed-reality scene.

With these adjustments in place, the patient-specific model could be imported and prepared automatically, without requiring manual intervention in the Unity Editor after each export from 3D Slicer. This automated pipeline provided a consistent and ready-to-use anatomical scene, forming the basis for the subsequent mixed-reality functionalities. After import and coordinate adaptation, the patient-specific anatomical scene could be visualized in the mixed-reality environment in different configurations, depending on the anatomical groups enabled for the procedure. Examples of the main model configurations visualized in mixed reality are shown in Figure 3.14.

These different anatomical configurations were then managed interactively

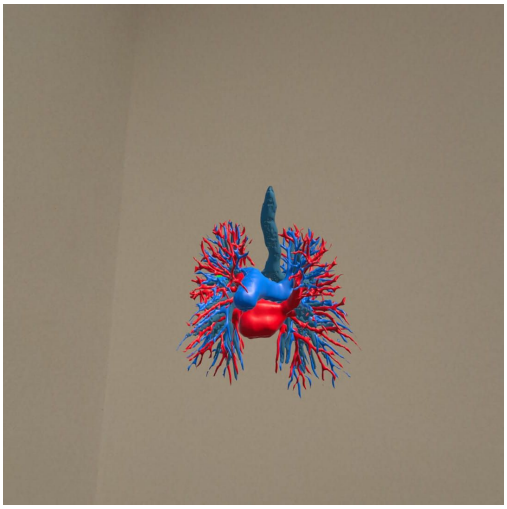

(a) AWVessels

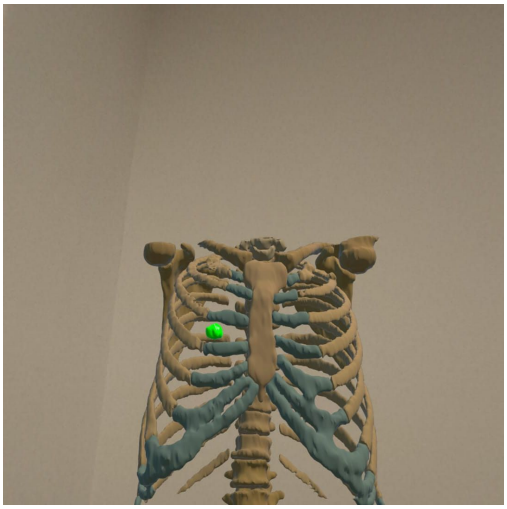

(b) Bones

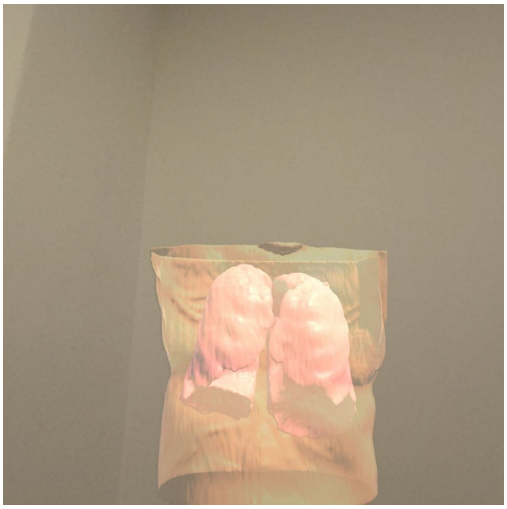

(c) Lungs and transparent skin

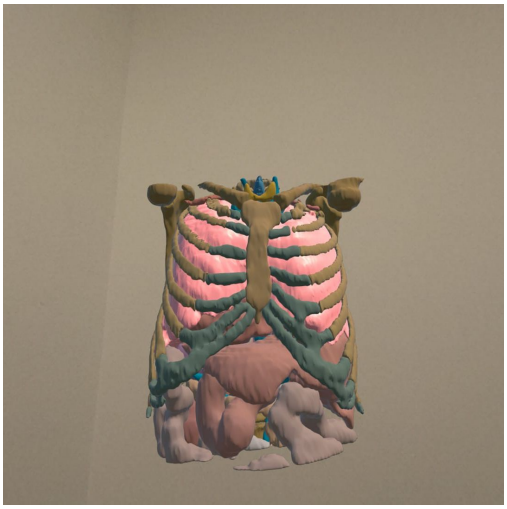

(d) Total

Figure 3.14: Examples of the patient-specific anatomical models visualized in mixed reality in Unity on Meta Quest 3. Different anatomical configurations could be displayed according to the visualization needs, including airways and vessels (AWVessels), bones, lungs with transparent skin, and the complete thoracic model.

through the graphical user interface, which allowed the user to selectively display or hide the available anatomical groups.

3.3.2 Graphical user interface and anatomical model management

After the patient-specific model was imported, the Unity application provided a graphical user interface to manage the anatomical scene and the main mixed-reality functions. The interface was designed to allow the operator to interact with the model directly inside the Meta Quest 3 environment, without requiring manual modifications in the Unity Editor after each export from 3D Slicer. The complete interface is shown in Figure 3.15.

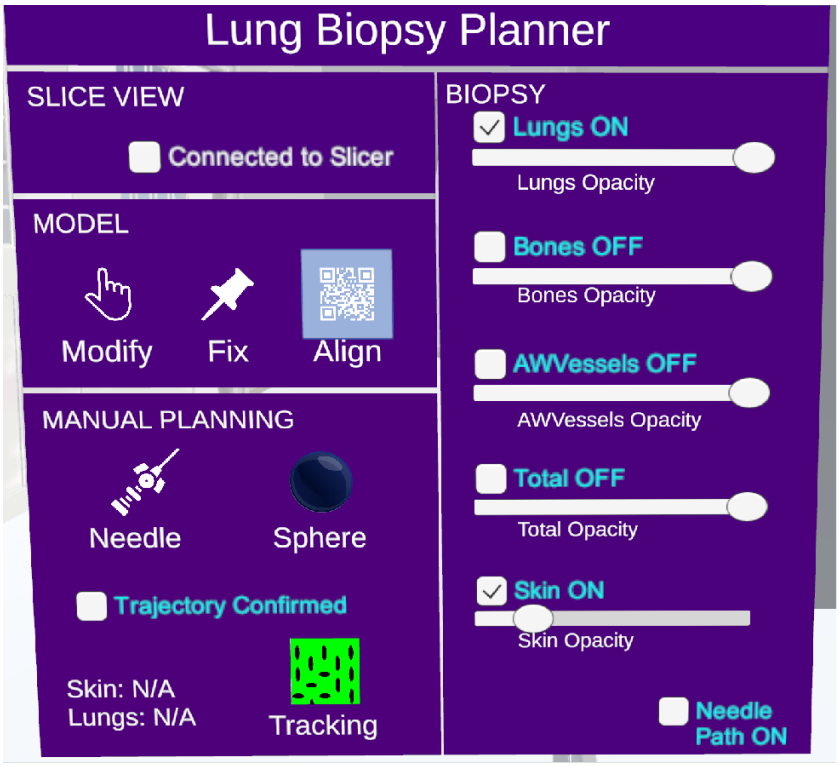

Figure 3.15: Graphical user interface developed in Unity for the Meta Quest 3 application. The interface provides controls for OpenIGTLink-based slice visualization, model manipulation, QR-based registration, manual trajectory planning, anatomical structure visibility, opacity regulation, and needle tracking.

The imported thoracic scene was organized into anatomical groups according to the structures generated during the pre-operative planning workflow. These groups

included the lungs, pulmonary nodule, bones, vessels and airways, skin surface, and planned instrument representation. Each group could be managed independently through the interface, allowing the user to show or hide specific structures and adjust their opacity according to the information required during visualization and guidance.

The graphical interface was divided into four main functional areas:

- **Slice view.** The *Connected to Slicer* control was used to activate the connection with 3D Slicer through OpenIGTLink. When the connection was enabled, a movable slicing plane appeared within the mixed-reality scene and could be used to cut through the imported anatomical model. The spatial position of this plane was sent from Unity to 3D Slicer, where the corresponding CT reslice was generated. The resulting slice was then transmitted back to the mixed-reality application and displayed on a virtual screen positioned next to the graphical interface. In this way, the operator could move the cutting plane through the 3D model and simultaneously visualize the corresponding CT section on the adjacent screen.

- **Model.** The *Modify* control enabled direct manipulation of the imported patient-specific model, allowing the user to translate, rotate, and scale it within the mixed-reality environment. The *Fix* control disabled model manipulation, preventing accidental changes in position or orientation once the model had been placed correctly. The *Align* control activated the QR-based registration procedure, allowing the virtual model to be aligned with the physical QR-code reference frame.

- **Manual planning.** The *Needle* control activated the virtual laser guidance tool used for manual trajectory selection. The *Sphere* control allowed the grabbable sphere used to orient the laser to be hidden when necessary, reducing visual occlusion of the anatomical model. When a valid trajectory was identified, the *Trajectory Confirmed* control enabled the visualization of a red marker on the skin surface, representing the selected skin entry point. The numerical fields associated with *Skin* and *Lungs* displayed the distance between the target nodule and, respectively, the skin entry point and the pleural surface, providing an estimate of the overall insertion depth and of the intrapulmonary path. The *Tracking* control enabled the camera-based tracking of the real needle.

- **Biopsy.** The *Biopsy* panel provided selective control of the anatomical structures included in the imported scene. The user could enable or disable the lungs, bones, airways and vessels, the general anatomical group obtained from the `total` segmentation task, and the skin surface. Each anatomical

group was associated with an opacity slider, allowing transparent visualization of external or overlapping structures. Finally, the *Needle Path ON* control enabled the visualization of the planned instrument representation generated in 3D Slicer through the PortPlacement module.

Overall, the graphical user interface represented the main interaction layer between the user and the patient-specific model. It allowed the anatomical scene generated in 3D Slicer to be explored in a flexible way and provided the controls required to activate the subsequent mixed-reality guidance functions, including registration, dynamic CT slice visualization, manual trajectory planning, and real needle tracking.

3.3.3 QR-based automatic registration

After importing the patient-specific model, an automatic registration procedure was implemented to align the virtual anatomy with the physical reference frame. The registration was based on four QR codes arranged at the corners of a square reference pattern. Each QR code contained a payload identifying its position within the square, namely lower-left, lower-right, upper-left, or upper-right. This allowed the Unity application to recognize not only the pose of each detected marker, but also its corresponding location within the predefined reference geometry.

The system was designed so that the origin of the 3D model coincides exactly with the geometric center of the square formed by the four QR codes. This point is not merely an arbitrary reference but corresponds to the isocenter used during the CT scan. Consequently, the exported anatomical model is already intrinsically aligned with this reference system.

Since the configuration of the QR codes remains rigidly fixed relative to the patient, the position of their center does not change between the scanning phase and the mixed-reality visualization phase. This ensures that the same physical point represents both the origin of the coordinate system during the CT scan and the spatial reference detected by the headset.

Registration is therefore based on the coincidence of these two elements: on one hand, the origin of the 3D model; on the other, the center of the QR marker system detected in real time. When the Meta Quest 3 identifies the QR codes, it is able to estimate the pose of the reference square in space. Since the model was constructed with the same origin, it is sufficient to apply this transformation to automatically achieve the correct alignment, as shown in Figure 3.16.

During the procedure, the model's position and orientation are continuously updated based on the placement of the QR code squares. This strategy ensures robust registration even when the reference system moves. If the QR codes are placed directly on the patient or on a structure rigidly attached to the patient, the

virtual model will automatically follow those movements, maintaining alignment with the actual anatomy.

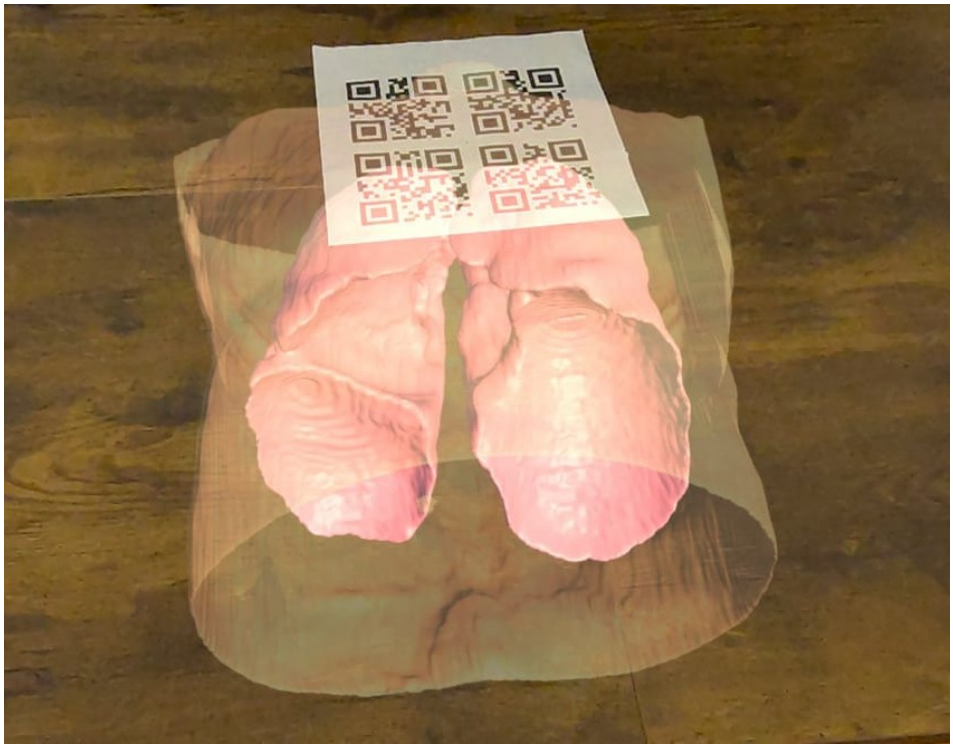

Figure 3.16: Example of QR-based automatic registration. The patient-specific anatomical model is aligned in MR with the four-QR-code reference system.

3.3.4 CT slice visualization through OpenIGTLink

In addition to the static anatomical model imported from the GLB file, the Unity application was designed to visualize dynamic CT slices within the mixed-reality scene. This functionality was implemented to allow the operator to inspect the internal volumetric information directly in relation to the patient-specific 3D model, without transferring the complete CT volume to the headset.

The communication between Unity and 3D Slicer was based on the OpenIGTLink protocol, which supports the exchange of image and transform data between image-guided therapy applications [56]. In the proposed workflow, the OpenIGTLink server was first activated from 3D Slicer. The Unity application then connected to this server as a client, enabling the exchange of spatial information and image data between the two environments.

After the connection was established, the Volume Reslice Driver module in 3D Slicer was configured to generate the axial CT section corresponding to the position of a slicing plane controlled in Unity [59]. When the slice visualization function was enabled in the mixed-reality application, a movable cutting plane appeared inside the imported anatomical model. This plane intersected the 3D model and represented the spatial position at which the corresponding CT slice had to be extracted.

The slicing plane could be translated along the model using an associated grabbable cube. By manipulating this cube, the user could move the plane through the patient-specific anatomy and explore different depths of the CT volume. The updated position of the plane was sent from Unity to 3D Slicer through OpenIGTLink. 3D Slicer then generated the corresponding CT reslice and transmitted it back to Unity.

The received slice was displayed on a virtual screen positioned next to the graphical user interface, as shown in Figure 3.17. In this way, the operator could move the cutting plane through the 3D anatomical model and simultaneously visualize the corresponding axial CT section in the mixed-reality environment.

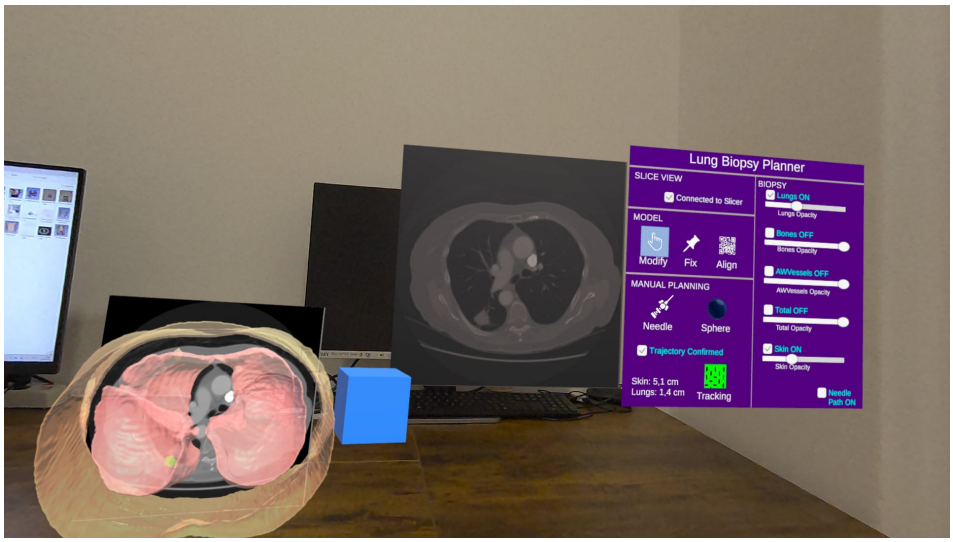

Figure 3.17: Dynamic CT slice visualization in the mixed-reality application. A movable slicing plane intersects the imported anatomical model, while the corresponding CT section is displayed on a virtual screen next to the graphical user interface.

This approach separated the static and dynamic components of the mixed-reality visualization. The GLB file provided the patient-specific 3D anatomical scene, while OpenIGTLink enabled the synchronization of the movable slicing plane with the corresponding CT section generated in 3D Slicer. In this way,

the application combined surface-based anatomical models with sectional image information, supporting a more complete inspection of the planned biopsy trajectory in the mixed-reality environment.

3.3.5 Virtual laser guidance for manual trajectory selection

A virtual laser guidance tool was implemented to allow manual exploration and selection of the biopsy trajectory directly within the mixed-reality scene. Although automatic trajectory analysis could support the identification of feasible access regions, the final selection of a percutaneous path was kept under user control. This choice was motivated by the need to preserve clinical interpretability and allowed the operator to evaluate the trajectory according to anatomical, procedural, and practical considerations that may not be fully captured by an automatic planning algorithm.

The tool was based on an interactive, grabbable sphere positioned within the mixed-reality scene. The user could manipulate the sphere manually and orient it toward the anatomical model. A virtual laser beam was emitted from the sphere, visually representing the candidate needle direction. In addition, a virtual crosshair aligned with the laser beam was displayed and dynamically positioned on the skin surface. The crosshair continuously adjusted its orientation to remain perpendicular to the local skin surface, providing a stable visual reference to accurately aim at the target nodule during trajectory selection. In this way, the operator could explore different access paths by moving and rotating the sphere, while directly observing the spatial relationship between the laser beam, the target nodule, and the surrounding anatomical structures (Figure 3.18).

The laser provided immediate visual feedback on the validity of the selected trajectory. By default, it was displayed in red and turned green only when the ray reached the target nodule without intersecting anatomical structures considered as obstacles, such as ribs, airways, and pulmonary vessels. This color-coded feedback allowed the operator to easily distinguish between obstructed trajectories and those that could potentially reach the target along a clear path (Figure 3.18b–c).

When a valid trajectory was identified, the graphical user interface displayed quantitative information related to the selected path. In particular, the system reported the distance between the nodule and the pleural surface, as well as the distance between the nodule and the skin entry point. These measurements provided an estimate of both the intrapulmonary path and the overall insertion depth, supporting the comparison of different candidate trajectories.

After visually evaluating the trajectory, the user could confirm the selection. Once confirmed, the application generated a highlighted red point on the skin surface corresponding to the chosen entry point. This point represented the planned puncture location associated with the selected trajectory and served as a spatial

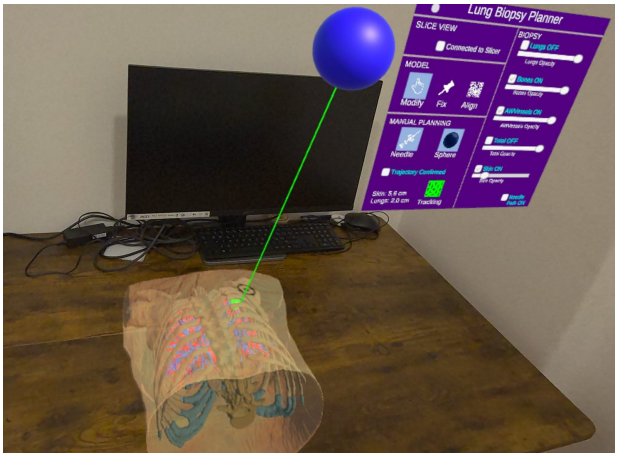

(a)

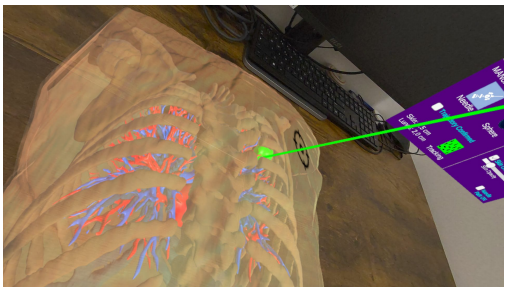

(b)

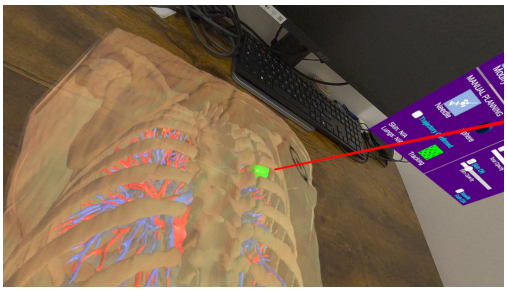

(c)

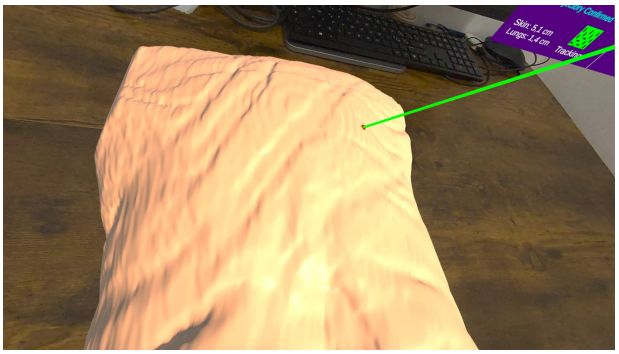

(d)

Figure 3.18: Virtual laser guidance for manual trajectory selection in the mixed-reality application. (a) General view of the grabbable sphere used to control the candidate trajectory. (b) Valid trajectory reaching the target nodule, shown in green. (c) Invalid trajectory intersecting the ribs, shown in red. (d) Skin entry point displayed after trajectory confirmation.

reference for subsequent guidance and instrument tracking (Figure 3.18d).

Overall, the virtual laser guidance tool was designed as an interactive planning aid rather than an autonomous decision-making system. It allowed the operator to define and confirm a trajectory manually, while receiving real-time visual feedback on target reachability, obstacle avoidance, and estimated insertion depth.

3.3.6 Camera-based tracking of the real needle

The final component of the Unity application was the camera-based tracking of the real needle. This functionality was implemented to estimate the pose of the instrument in the mixed-reality scene and to compare its orientation with the trajectory selected through the virtual laser guidance system. The tracking pipeline was based on the analysis of camera frames acquired from the Meta Quest 3 and processed by an external computer-vision server.

The image acquisition step relied on the Meta Passthrough Camera API, which provides access to the forward-facing RGB cameras of the headset for computer-vision applications [54]. In the implemented setup, the application had to be installed and executed directly on the Meta Quest 3 to access the camera stream used for tracking. The acquired frames were then sent from Unity to the Flask server, where the marker detection and pose estimation pipeline was executed.

The tracked instrument was equipped with a cylindrical marker inspired by the marker design proposed by Cartucho et al. for surgical tool tracking [60]. The marker consisted of a green background with black elliptical features arranged to encode a binary pattern around the cylindrical surface. The processing pipeline first segmented the green region in order to isolate the marker from the rest of the image. After this step, the internal black features were detected and used to extract the marker keypoints.

The detected keypoints were then analyzed to reconstruct the binary sequences encoded on the cylindrical marker. This step allowed the system to associate the two-dimensional keypoints detected in the image with the corresponding three-dimensional points defined on the known marker geometry. Once these correspondences were established, the pose of the marker with respect to the camera was estimated using a Perspective-n-Point approach with RANSAC. In OpenCV, this method estimates the rotation and translation of an object from known correspondences between three-dimensional object points and their two-dimensional image projections, given the camera intrinsic parameters [64].

A separate camera calibration procedure was not required in this implementation, since the Passthrough Camera API provided access to the intrinsic parameters of the Meta Quest 3 camera. The distortion coefficients were therefore not estimated through an external calibration procedure and were treated as null in the pose-estimation pipeline, relying on the image stream and camera parameters provided

by the headset software stack.

After pose estimation, the three coordinate axes of the marker were overlaid on the image as a visual check of the tracking result. The estimated marker pose was then converted into the Unity world coordinate system by combining the pose returned by the computer-vision pipeline with the pose of the headset camera at the time of image acquisition. This allowed the position and orientation of the real instrument to be updated in real time within the mixed-reality scene.

Once the needle pose was available in Unity, the application compared the tracked instrument with the manually planned trajectory defined through the virtual laser guidance system. In particular, the lateral deviation from the planned path and the angular deviation between the instrument axis and the selected trajectory were evaluated (Figure 3.19). These quantities were used to provide preliminary needle-guidance feedback, allowing the operator to understand how the real instrument should be translated and oriented to align with the planned path.

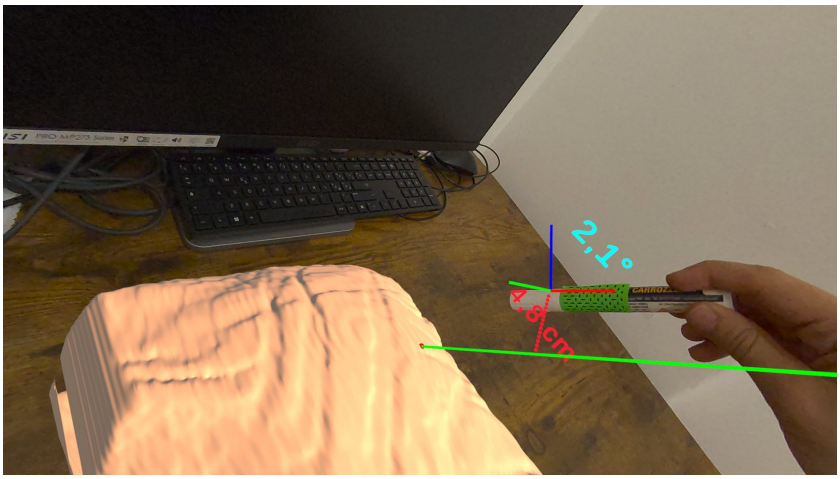

Figure 3.19: Camera-based needle tracking and preliminary guidance feedback. The tracked needle axis is compared with the manually planned trajectory, allowing the estimation of lateral deviation and angular deviation.

3.4 Experimental evaluation

A preliminary experimental evaluation was designed to assess the feasibility of the proposed mixed-reality workflow in a controlled setup. The evaluation was organized into two components. The first component consisted of a phantom-based test, aimed at verifying the practical use of the QR-code-based registration strategy and the visualization of the patient-specific model in relation to a physical

thoracic phantom. The second component consisted of a questionnaire administered to potential clinical users, aimed at collecting preliminary feedback on usability, anatomical visualization, and perceived clinical usefulness of the system.

The phantom-based experiment was not intended to provide a complete clinical validation of the guidance system or a quantitative assessment of targeting accuracy. Instead, it was designed as an initial feasibility test to verify whether the virtual model generated from the CT scan could be aligned with the physical phantom using the QR-code reference system and, when necessary, manually refined using an additional calibration volume.

3.4.1 Phantom setup

The experimental setup was based on a custom thoracic phantom designed to reproduce, in simplified form, the main anatomical and procedural elements relevant to CT-guided lung biopsy. The phantom consisted of an agar-agar gelatin matrix, used to mimic the soft-tissue volume, with five embedded artificial nodules made of modelling clay. In addition, 3D-printed bones were included to provide rigid anatomical structures and to reproduce the presence of obstacles that could influence trajectory planning and mixed-reality visualization (Figure 3.20).

Figure 3.20: Custom thoracic phantom used for the preliminary evaluation.

The phantom was positioned on a CT table and scanned using the same general workflow adopted for the patient-specific model generation. Before CT acquisition, a four-QR-code reference pattern was placed in a fixed position with respect to the phantom. This configuration was used to define the physical reference system for the subsequent mixed-reality registration. The QR-code arrangement was designed

so that its geometric center could be used as the spatial reference for aligning the virtual model with the physical setup.

An additional LEGO structure was positioned between the phantom and the QR-code reference pattern during the CT scan. This object was included as a calibration volume and served as a visual reference for manual refinement of the alignment. In the mixed-reality scene, the corresponding virtual LEGO model could be visualized together with the anatomical phantom model. Therefore, if the automatic QR-based registration did not produce a visually satisfactory alignment, the user could manually adjust the model position and orientation until the physical LEGO structure overlapped with its virtual counterpart. The phantom setup is shown in Figure 3.21.

Figure 3.21: Phantom setup used for the preliminary experimental evaluation. The phantom consisted of a gelatin matrix containing artificial nodules and 3D-printed bones. A LEGO calibration volume and a four-QR-code reference pattern were included to support mixed-reality registration and manual alignment refinement.

After CT acquisition, the phantom data were processed in 3D Slicer to generate the corresponding three-dimensional scene and subsequently loaded into the Unity

application, where it could be visualized in mixed reality and registered to the physical phantom through the QR-code reference system.

3.4.2 Clinical user questionnaire

In addition to the phantom-based test, a preliminary questionnaire was administered to potential clinical users after visualization and interaction with the mixed-reality application. The purpose of the questionnaire was to collect subjective feedback on the perceived usefulness, clarity, usability, workload, and potential clinical relevance of the proposed workflow. This evaluation was not intended as a formal usability validation, given the limited number of participants, but as an initial assessment of the system from the perspective of users familiar with image-guided procedures.

The evaluation included both a custom questionnaire, specifically designed for the developed mixed-reality workflow, and two standardized usability assessment tools: the System Usability Scale (SUS) and the NASA Task Load Index (NASA-TLX). The custom questionnaire was used to investigate system-specific aspects, such as anatomical visualization, model interaction, trajectory planning, needle tracking, clinical usefulness, and headset comfort. The SUS was included to obtain a general measure of perceived system usability, while the NASA-TLX was used to estimate the subjective workload associated with the interaction task [65, 66].

The custom questionnaire was based on a 5-point Likert scale, where 1 corresponded to *strongly disagree*, 2 to *disagree*, 3 to *neutral*, 4 to *agree*, and 5 to *strongly agree*. The items were grouped into five main domains, corresponding to the principal components of the developed system:

- **MR visualization and anatomical understanding**, aimed at evaluating the perceived quality of the patient-specific anatomical models and their usefulness for understanding the spatial relationship between the target nodule and surrounding structures;
- **MR interaction and interface**, aimed at assessing the clarity of the graphical user interface, the possibility of interacting with the model, and the usefulness of selective anatomical visualization;
- **Needle tracking and guidance feedback**, aimed at collecting feedback on the virtual laser guidance, the visualization of the planned trajectory, and the preliminary feedback provided by lateral and angular deviation;
- **Clinical usefulness**, aimed at evaluating whether the proposed workflow could be perceived as useful for supporting planning, anatomical interpretation, and future guidance during CT-guided lung biopsy;

- **Comfort and usability**, aimed at assessing the practical acceptability of the Meta Quest 3 headset and the overall usability of the system during a short simulation session.

Both positively and negatively worded statements were included in the custom questionnaire. The negatively worded items were introduced to reduce response bias and to verify the consistency of the answers. These items were considered separately during interpretation, so that high agreement with a negatively worded statement indicated a potential limitation of the system rather than a positive evaluation.

The SUS questionnaire consisted of ten standard statements rated on a 5-point Likert scale, alternating positively and negatively worded items. The SUS score was computed using the standard scoring procedure:

$$\text{SUS score} = 2.5 \left[\sum (Q_{\text{odd}} - 1) + \sum (5 - Q_{\text{even}}) \right]$$

where Q_{odd} represents the responses to positively worded items and Q_{even} represents the responses to negatively worded items. The resulting score ranges from 0 to 100, with higher values indicating better perceived usability.

The NASA-TLX questionnaire was used to assess the perceived workload associated with the use of the system. Six dimensions were considered: mental demand, physical demand, temporal demand, performance, effort, and frustration. Each dimension was first converted to a 0–100 rating scale. The final workload score was then computed as the arithmetic mean of the six dimension scores:

$$\text{NASA-TLX score} = \frac{\text{MD} + \text{PD} + \text{TD} + \text{P} + \text{E} + \text{F}}{6}$$

where MD, PD, TD, P, E, and F correspond respectively to mental demand, physical demand, temporal demand, performance, effort, and frustration. Lower NASA-TLX scores indicated lower perceived workload during interaction with the mixed-reality application.

Overall, the combination of a custom questionnaire with SUS and NASA-TLX allowed the evaluation to capture both system-specific clinical feedback and standardized usability-related measures.

Chapter 4

Results

4.1 Phantom-based evaluation of QR-based alignment

The first experimental result concerned the preliminary phantom-based evaluation of the QR-code-based registration workflow. The test was performed using the custom thoracic phantom described before, with the aim of verifying whether the virtual model generated from the CT acquisition could be visualized and aligned with the physical phantom in the mixed-reality environment.

After CT acquisition, the phantom volume was processed in 3D Slicer to generate the corresponding virtual model. In this case, the automatic segmentation performed with TotalSegmentator did not provide usable anatomical results. This was expected to some extent, since the phantom did not reproduce the intensity distribution and anatomical complexity of a real thoracic CT scan. For this reason, manual segmentation was performed instead. The manually segmented structures included the external surface of the phantom, the artificial nodules, and the LEGO calibration volume. However, it was not possible to reliably separate the 3D-printed bones from the surrounding phantom material. Consequently, the mixed-reality model used for the alignment test included only the phantom surface, the nodules, and the calibration object.

The QR-code-based registration procedure was then tested in the mixed-reality application. After detection of the four QR codes by the Meta Quest 3, the virtual model was automatically positioned with respect to the QR-code reference system. The initial automatic alignment was successful in terms of model visualization and spatial registration activation, but it showed a vertical offset of approximately 9 cm between the virtual model and the physical phantom.

This offset was corrected manually using the LEGO calibration volume as a visual reference. Since this structure had been included in the CT scan and was also

present in the physical setup, its virtual counterpart could be superimposed onto the real object observed through the headset. By manually adjusting the position of the virtual model, the LEGO volumes were successfully aligned, resulting in a satisfactory visual correspondence between the virtual phantom model and the physical setup. Figure 4.1 shows the mixed-reality visualization before and after manual refinement of the alignment in this test.

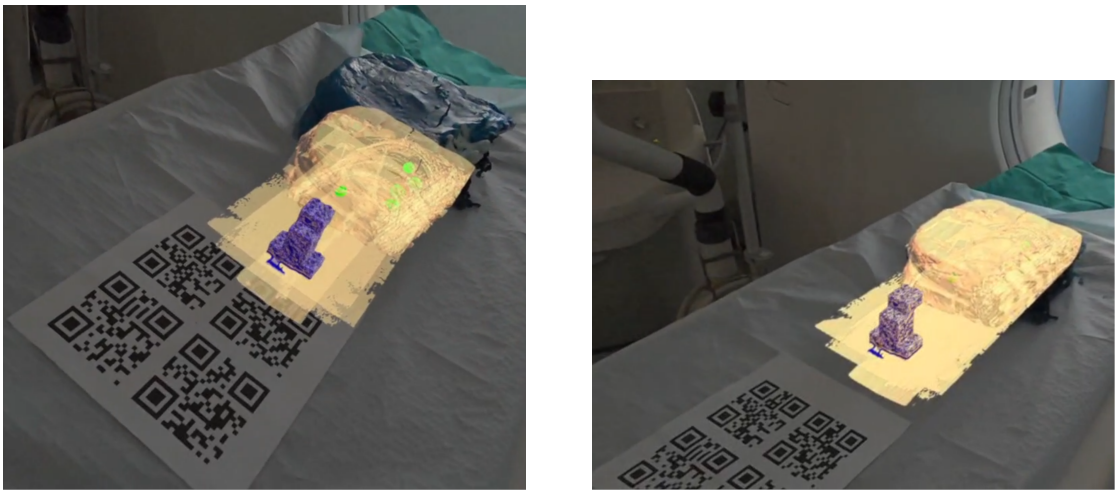

(a) Before manual refinement. **(b)** After manual refinement.

Figure 4.1: Mixed-reality visualization before and after manual alignment refinement using the calibration volume as a spatial reference.

Overall, the phantom test demonstrated the feasibility of the proposed QR-based registration workflow in a controlled experimental setup. Although the automatic alignment required manual refinement, the presence of a simple calibration volume allowed the offset to be corrected in a practical and intuitive way. From a procedural perspective, placing a small calibration object between the patient or phantom and the four-QR-code reference pattern appears compatible with a possible future clinical setup, provided that the object does not interfere with the sterile field, patient access, or CT acquisition.

A limitation observed during the test was the limited precision of manual calibration when the model was grabbed and moved directly using the Meta Quest controllers. This interaction modality was sufficient for a coarse correction, but it was not optimal for fine alignment. Based on feedback from the radiologists involved in the evaluation, a more precise calibration strategy should be implemented in future versions of the application. In particular, model adjustment could be controlled through controller buttons or joystick-based incremental movements, allowing the virtual model to be translated by predefined fixed steps. This would

make the manual refinement more reproducible and less dependent on the stability of freehand controller manipulation.

Therefore, the phantom-based evaluation confirmed the preliminary technical feasibility of QR-code-based model alignment, while also identifying the need for a more accurate and controllable manual calibration interface.

4.2 Clinical user questionnaire results

The second component of the preliminary evaluation consisted of a questionnaire battery administered to two physicians, one of whom was an interventional radiologist. The questionnaires were completed after interaction with the mixed-reality application and were designed to collect preliminary feedback on anatomical visualization, interaction, trajectory guidance, needle tracking, clinical usefulness, comfort, perceived usability, and workload during use.

The evaluation included a custom questionnaire based on a 5-point Likert scale, the SUS, and the NASA-TLX, as described in the Materials and Methods chapter. The custom questionnaire was used to assess system-specific aspects of the proposed workflow, whereas SUS and NASA-TLX were used to provide standardized usability and workload measures. For each questionnaire, only descriptive results were reported.

4.2.1 Custom questionnaire results

The custom questionnaire was based on a 5-point Likert scale, where 1 corresponded to *strongly disagree*, 2 to *disagree*, 3 to *neutral*, 4 to *agree*, and 5 to *strongly agree*. The items were grouped into five domains: MR visualization and anatomical understanding, MR interaction, needle tracking and guidance feedback, clinical usefulness, and comfort and usability. For each item, the mean score and standard deviation were computed from the available responses and are reported in Tables 4.1 and 4.2.

Table 4.1: Clinical user questionnaire results regarding system features and clinical applicability. Scores are reported as mean values ± standard deviation on a 5-point Likert scale. Negatively worded items are marked with an asterisk (*).

Evaluated statement	**Mean score ± SD**
MR visualization and anatomical understanding	
The visual quality and level of detail of the mixed-reality 3D models, including lungs and nodules, were high.	4.50 ± 0.71
The 3D holographic model improved my spatial understanding of the target nodule relative to surrounding anatomical structures.	5.00 ± 0.00
Visualizing the internal CT slice directly within the 3D model provided clear clinical utility during the planning phase.	4.50 ± 0.71
The mixed-reality 3D models were difficult to interpret because of insufficient visual clarity or detail.*	3.00 ± 1.41
MR interaction	
The mixed-reality user interface was easy to understand during the simulation.	4.50 ± 0.71
The manipulation of the 3D model was sufficiently precise and responsive for adjusting its position, orientation, and scale during the simulation.	3.00 ± 1.41
The laser-guided trajectory helped me reach the target nodule more easily.	4.50 ± 0.71
Needle tracking and guidance feedback	
The needle tracking appeared sufficiently stable for a prototype demonstration.	4.00 ± 1.41
The displayed lateral and angular deviations helped me to understand how to align the instrument with the planned trajectory.	3.00 ± 1.41
The needle tracking appeared unstable or unreliable during the prototype demonstration.*	3.50 ± 0.71
Clinical usefulness	
The system offered a significant improvement in understanding the optimal needle path compared to traditional methods.	4.00 ± 0.00
Overall, I consider this system a promising tool for supporting lung biopsy planning and guidance.	5.00 ± 0.00
The system did not provide a clear advantage over traditional visualization methods for understanding the optimal needle path.*	1.50 ± 0.71

Table 4.2: Clinical user questionnaire results regarding system comfort and usability. Scores are reported as mean values ± standard deviation on a 5-point Likert scale. Negatively worded items are marked with an asterisk (*).

Evaluated statement	**Mean score ± SD**
Comfort and usability	
I experienced noticeable visual fatigue during or immediately after the session.*	2.00 ± 1.41
The Meta Quest 3 headset was comfortable enough for the duration of the session.	3.50 ± 0.71
I experienced minimal visual fatigue during or immediately after the session.	4.00 ± 1.41
Overall, the system seemed too complex for practical use in a clinical workflow.*	3.00 ± 1.41

4.2.2 System Usability Scale results

The SUS questionnaire was used to obtain a standardized measure of perceived system usability. The mean SUS score was 67.50/100. This value was slightly below the commonly used usability reference threshold of 68/100, while remaining close to it.

4.2.3 NASA-TLX results

Table 4.3: NASA-TLX workload results. The final score corresponds to the mean NASA-TLX score.

Workload dimension	**Mean score**
Mental demand	75.00
Physical demand	42.50
Temporal demand	40.00
Performance	45.00
Effort	62.50
Frustration	37.50
Overall NASA-TLX score	**50.42**

The NASA-TLX questionnaire was used to assess the perceived workload associated with the interaction with the mixed-reality application. The six workload

dimensions were reported on a 0–100 scale and averaged across the available responses. The final NASA-TLX score was also computed according to the procedure described in the previous chapter.

Chapter 5

Discussion

5.1 Interpretation of the preliminary results

The preliminary evaluation of the proposed system provided two complementary types of information. The phantom-based test assessed the technical feasibility of the QR-code-based alignment strategy in a controlled setup, while the clinical user questionnaire provided initial feedback on the perceived usefulness, usability, and limitations of the mixed-reality workflow. Since the evaluation involved a limited number of participants and a single phantom setup, the results should not be interpreted as a clinical validation, but rather as an exploratory assessment of the developed prototype.

5.1.1 Interpretation of the phantom-based alignment test

The phantom-based evaluation provided an important indication of the potential and current limitations of the QR-code-based registration strategy. After detection of the four QR codes, the virtual model was automatically positioned in the mixed-reality scene, confirming that the reference pattern could be used to initialize the alignment between the CT-derived model and the physical setup. However, the initial registration showed a vertical offset of approximately 9 cm between the virtual and physical phantom.

Since the misalignment was mainly observed along the depth direction, this offset was likely related to an inaccurate estimation of the QR-code distance from the headset cameras. In other words, the headset was able to correctly recognize the relative positions of the QR markers on the reference plane, and therefore to reconstruct the square pattern. However, the estimation of the distance between the QR-code plane and the headset introduced a residual displacement of the virtual model along the depth direction.

Despite this initial offset, the use of the LEGO calibration volume allowed the alignment to be manually refined in an effective way. This provided a simple and intuitive visual reference for correcting the residual registration error and obtaining a satisfactory correspondence between the mixed-reality model and the phantom.

This result suggests that the QR code-based registration strategy has great potential as a method for automatically aligning the virtual model with the patient. In a possible clinical scenario, the calibration object should be positioned outside the patient but included in the CT scan, without, however, interfering with the sterile field or access to the patient.

However, the phantom test also highlighted that the current manual manipulation strategy was not sufficiently precise for optimal calibration. Grabbing and moving the model with the Meta Quest controllers allowed a coarse correction, but it did not provide the level of control required for fine alignment.

Overall, the phantom-based test confirmed the feasibility of the proposed registration workflow. Although the automatic alignment was not perfect in the tested configuration, the combination of QR-code-based registration and calibration-volume-based manual refinement provided an effective alignment strategy. These results support the potential of the system for automatic model registration in a future patient-based workflow.

5.1.2 Interpretation of questionnaire scores

The questionnaire results provided a preliminary indication of how the proposed mixed-reality system was perceived by clinical users. Overall, the custom questionnaire showed that the most appreciated aspects were related to anatomical visualization, spatial understanding, and trajectory planning. The visualization domain obtained high scores, with the improvement in spatial understanding of the target nodule relative to surrounding anatomical structures reaching 5.00 ± 0.00. Similarly, both the visual quality of the 3D models and the integration of the internal CT slice within the mixed-reality scene received positive evaluations, suggesting that the system was considered useful for interpreting the anatomical context during the planning phase.

These findings are particularly relevant for the intended application of the system. One of the main limitations of conventional CT-guided lung biopsy is the need to mentally reconstruct a three-dimensional anatomical configuration from two-dimensional CT slices. The positive feedback obtained for anatomical visualization suggests that the proposed mixed-reality model may help reduce this cognitive translation step by allowing the operator to inspect the target nodule, surrounding structures, and candidate trajectory in a more intuitive spatial format.

Trajectory planning was also perceived positively. The virtual laser guidance tool received a mean score of 4.50 ± 0.71, indicating that the possibility of manually

exploring candidate access directions directly in mixed reality was considered useful. This supports one of the main design choices of the system: preserving user control over trajectory selection while providing immediate visual feedback within the 3D scene.

The clinical usefulness domain confirmed this positive trend. The system was perceived as promising for supporting lung biopsy planning and guidance, with the corresponding statement receiving the maximum score of 5.00 ± 0.00. In addition, the low agreement with the negatively worded statement regarding the absence of clear advantages over traditional visualization methods suggests that the users recognized a potential added value of the proposed workflow for anatomical understanding and trajectory planning.

From the usability perspective, the interface was generally considered easy to understand, while headset comfort and visual fatigue showed acceptable but more variable results. This suggests that the short interaction session was overall tolerated, although comfort and fatigue should be further assessed in larger samples and during longer usage sessions.

The main limitations emerging from the custom questionnaire concerned model manipulation and needle tracking. Model manipulation received an intermediate score of 3.00 ± 1.41, indicating that direct grabbing of the model with the Meta Quest controllers or hands was not sufficiently precise for fine alignment. This result is consistent with the phantom-based test, where manual calibration was feasible but would benefit from more controlled adjustment tools, such as incremental translations and rotations.

Needle tracking was also evaluated more cautiously. Although the concept of comparing the real instrument with the planned trajectory was considered potentially useful, the scores related to lateral and angular deviation feedback and tracking reliability indicated that the current implementation did not yet provide the robustness and accuracy required for reliable procedural guidance. Therefore, this component should be interpreted as a preliminary proof of concept rather than as a validated navigation tool.

The standardized usability questionnaires provided additional information. The mean SUS score was 67.50/100, which is slightly below the commonly used acceptability threshold of 68/100. This result can be interpreted as borderline usability: the system was perceived as close to acceptable, but not yet fully mature from an interaction and workflow integration perspective. This is coherent with the prototype nature of the application and with the limitations identified in model manipulation and needle tracking.

The NASA-TLX results showed an overall workload score of 50.42/100. Among the individual workload dimensions, the highest values were observed for mental demand and effort, with mean scores of 75.00/100 and 62.50/100, respectively.

This suggests that interacting with the prototype still required a relevant cognitive effort, probably related to the interpretation of the 3D environment, the use of the interface, and the need to manage several mixed-reality functionalities. In contrast, temporal demand and frustration were lower, with mean scores of 40.00/100 and 37.50/100, respectively. This indicates that the participants did not perceive strong time pressure and did not find the experience particularly frustrating.

Overall, the questionnaire results suggest that the system was perceived as promising, especially for patient-specific 3D visualization, spatial understanding, and laser-based trajectory planning. At the same time, the SUS and NASA-TLX results confirm that the prototype still requires technical and interaction improvements before it can be considered mature for a clinical workflow. In particular, future development should focus on reducing cognitive workload, improving the precision of model manipulation, simplifying interaction, and increasing the robustness of needle tracking.

5.2 Main limitations of the study

Although the preliminary results confirmed the technical feasibility of the proposed mixed-reality workflow, several limitations must be considered. These limitations concern both the experimental evaluation and the current implementation of the system, and should be addressed before any possible clinical translation.

5.2.1 Restricted Validation and Lack of Clinical Evidence

A first limitation is related to the preliminary nature of the evaluation. The system was tested on a single custom phantom and the clinical feedback was collected from only two physicians. Therefore, the results cannot be interpreted as a validation of clinical accuracy, safety, or procedural effectiveness. The phantom test provided useful information on model registration and calibration, while the questionnaire allowed an initial assessment of perceived usability and usefulness. However, larger studies involving more users, multiple phantom configurations, and eventually ex-vivo or clinical scenarios would be required to assess the real impact of the system on biopsy planning and needle guidance.

5.2.2 Limited Anatomical Realism of the Phantom Model

A second limitation concerns the simplified anatomical realism of the phantom. The phantom reproduced some relevant elements of the biopsy scenario, including the soft-tissue volume, artificial nodules, and bones. However, it did not reproduce the complexity of real thoracic anatomy, respiratory motion, lung deformation,

or tissue displacement during needle insertion. Moreover, TotalSegmentator was not able to provide reliable automatic segmentation on the phantom CT dataset, probably because the phantom did not reproduce the intensity distribution of real anatomical structures. As a result, manual segmentation was required and only the external surface, nodules, and calibration volume could be separated, while the 3D-printed bones could not be reliably distinguished from the surrounding material. This limited the anatomical completeness of the model used during the phantom-based test.

5.2.3 Registration accuracy and manual alignment refinement

Another important limitation is related to the overall accuracy of the registration process. In the proposed workflow, model alignment was obtained through the combination of QR-code-based automatic registration and manual refinement using the calibration volume. The QR-code-based procedure successfully initialized the alignment between the virtual model and the physical setup, but the initial result showed a vertical offset of approximately 9 cm. Although the offset could be corrected using the calibration volume, the refinement step was limited by the current interaction modality. In the implemented application, the virtual model could be grabbed and moved using the Meta Quest controllers or hand interaction. This approach was sufficient for rough alignment, but it did not provide the sensitivity required for fine calibration. Small hand movements could produce excessive or poorly controlled changes in the model position, making precise alignment more difficult.

5.2.4 Limited Accuracy and Robustness of Needle Tracking

Needle tracking represents another relevant limitation of the current prototype. The implemented tracking system allowed the real instrument to be detected and compared with the planned trajectory, providing lateral and angular deviation feedback. However, based on the feedback provided by the interventional radiologists, its accuracy and stability were not yet sufficient for reliable clinical needle guidance. For this reason, the tracking should be considered a preliminary proof of concept rather than a validated navigation component.

At the same time, the concept of tracking-based feedback was considered useful by the other surgeon involved in the evaluation, particularly as a way to visualize the relationship between the real instrument and the planned trajectory. This suggests that, although the current implementation requires substantial improvement in terms of accuracy and robustness, the underlying guidance concept remains promising.

5.2.5 Computational Constraints of Standalone Implementation

A further technical limitation is related to the computational constraints of the standalone Meta Quest 3 implementation. In the current workflow, needle tracking required the application to run directly on the headset rather than being connected to the PC. This means that the application relied on the onboard computational resources of the Meta Quest 3, which are more limited than those of a dedicated workstation or desktop GPU. In this thesis, development and testing on PC could benefit from a more powerful graphics card, such as an NVIDIA RTX 3070, whereas the standalone headset required a lower rendering quality to maintain acceptable performance and avoid lag. This trade-off affected the visual quality and complexity of the scene that could be rendered while maintaining real-time interaction and tracking.

5.2.6 Lack of Validation in Real Clinical Environment

Finally, the system was not tested in a real procedural environment. Aspects such as sterility, patient positioning, respiratory motion, CT room constraints, lighting conditions, occlusion of QR codes or tracking markers, and interaction with clinical staff were not evaluated. These factors may significantly influence the usability and reliability of the system during an actual biopsy workflow. Therefore, although the proposed system showed promising results in a controlled phantom setup, further development and validation are required before it can be considered suitable for clinical use.

Overall, the main limitations of this work concern the preliminary scale of the evaluation, the simplified phantom setup, the residual registration error, the limited precision of manual model manipulation, the insufficient accuracy of the needle tracking component, and the computational constraints of the standalone headset.

5.3 Future developments

The limitations identified during the preliminary evaluation define several directions for future development. These improvements concern both the technical implementation of the system and the experimental validation required before any possible clinical translation.

- **More precise model manipulation and calibration.** One of the first improvements should concern the manual refinement of model alignment. In

the current implementation, the virtual model can be grabbed and moved using the Meta Quest controllers or hand interaction. Although this approach is intuitive, it does not provide sufficient precision for fine calibration. Future versions of the application should therefore include dedicated alignment controls, such as arrows integrated into the user interface or incremental movements controlled through the controller buttons or joysticks. Each input could translate or rotate the model by a predefined amount, allowing a more precise and reproducible refinement of the alignment with the calibration volume.

- **Larger and more realistic validation studies.** The experimental evaluation should also be expanded. The present work was limited to a single phantom setup and a small number of clinical users. Future studies should involve a larger group of physicians, including interventional radiologists with different levels of experience, in order to obtain more representative feedback on usability, clinical usefulness, and workflow integration. In parallel, the system should be tested on more complex and anatomically realistic phantoms, progressively including more realistic tissue densities, anatomical obstacles, target depths, and access trajectories. This would represent an intermediate step toward pre-clinical studies and, eventually, clinical validation.

- **Respiratory motion compensation.** A major future development concerns the integration of respiratory motion analysis. In the current prototype, the patient-specific model is treated as a static object, which is appropriate for a preliminary phantom-based evaluation but does not reproduce the real conditions of lung biopsy. In clinical practice, the position of the lung, target nodule, pleura, and surrounding structures changes during respiration. Future versions of the system should therefore account for the respiratory phase, either by acquiring the reference CT scan in a controlled breathing condition or by integrating respiratory tracking strategies.

  At a preliminary level, the respiratory excursion of the patient could be measured and displayed directly within the mixed-reality interface. These values could then be compared with the respiratory condition recorded during the preliminary CT acquisition used to generate the patient-specific model. In this way, the operator could verify whether the patient's current breathing state is comparable to the one represented by the virtual model, reducing the risk of using anatomical information that no longer corresponds to the real position of the target and surrounding structures.

  More advanced developments could include respiratory gating or dynamic model adaptation, allowing the virtual anatomy to be updated according to the estimated respiratory phase. This would improve the reliability of the

mixed-reality overlay and could reduce errors caused by breathing-induced displacement during needle planning and guidance.

- **Tissue deformation modelling.** In addition to respiratory motion, tissue deformation should be considered. During needle insertion, the interaction between the instrument and soft tissues may alter the position of the target and surrounding structures. This is particularly relevant in lung interventions, where tissue compliance, pleural displacement, gravity-induced deformation, and needle-induced deformation can affect the correspondence between the pre-procedural model and the real anatomy. Future developments could explore deformable models or biomechanical simulations to better estimate these changes and improve the reliability of the guidance information during needle advancement.

- **Improved needle tracking.** The needle tracking component should be further developed and quantitatively validated. In the current prototype, tracking was implemented as a preliminary proof of concept and did not yet reach the accuracy required for reliable procedural guidance. Future work should investigate more robust tracking solutions, including higher-resolution cameras, external camera systems, and improved marker detection algorithms. The tracking accuracy should also be evaluated against an external reference system, in order to quantify lateral and angular errors with respect to the planned trajectory.

- **Dedicated marker support for clinical instruments.** Another important improvement concerns the physical integration of the tracking marker with the interventional instrument. In the current implementation, the tracking concept was tested preliminarily using a marker attached to a pen-like object. In future versions, a dedicated 3D-printed support should be developed to hold the markers in a stable and reproducible configuration. This support could be designed as a clip attached to the cannula in the case of a coaxial biopsy technique, where the introducer remains in place during multiple sampling steps. Alternatively, in procedures performed with a single needle, the support could be attached directly to the needle or needle holder. This would make the tracking setup more realistic, stable, and compatible with the clinical workflow.

- **Voice-command interaction.** Finally, voice commands could be integrated to improve interaction with the application. During an interventional procedure, the operator may need to keep the hands available for patient preparation, instrument handling, or sterile tasks. Voice interaction could allow faster activation of key functions, such as changing model visibility,

adjusting opacity, enabling the laser, confirming the trajectory, activating tracking, or switching between visualization modes. This would make the interface more efficient and could reduce the need for repeated manual interaction with controllers or virtual panels.

Overall, future developments should focus on improving calibration precision, increasing tracking accuracy, modelling respiratory motion and tissue deformation, and validating the workflow in progressively more realistic scenarios. These steps are necessary to move from a preliminary technical feasibility prototype toward a robust mixed-reality platform for supporting CT-guided lung biopsy planning and guidance.

Chapter 6

Conclusions

This thesis presented the development and preliminary evaluation of an integrated mixed-reality workflow for supporting the planning and guidance of needle insertion in thoracic biopsy using Meta Quest 3. The work was motivated by one of the main limitations of conventional CT-guided lung biopsy: the need for the operator to mentally translate two-dimensional CT images into a three-dimensional understanding of the target lesion, surrounding anatomical structures, and planned needle trajectory.

The first contribution of the thesis was the integration of the pre-procedural workflow within 3D Slicer. A custom module was developed to organize the main planning steps in a single environment, including CT loading, anatomical segmentation, target definition, obstacle model generation, trajectory analysis, visual verification, OpenIGTLink configuration, and export of the patient-specific model. This integration reduced workflow fragmentation and provided a structured pipeline from medical image processing to mixed-reality visualization.

The second contribution was the development of a Unity-based mixed-reality application for Meta Quest 3. The application allowed the patient-specific thoracic model generated in 3D Slicer to be imported, visualized, and manipulated in mixed reality. It also included selective anatomical visualization, opacity control, QR-code-based registration, and dynamic CT slice visualization through OpenIGTLink. These functionalities allowed three-dimensional anatomical information and sectional imaging data to be inspected within the same mixed-reality environment.

The system also included preliminary guidance functionalities. A virtual laser tool was implemented to support manual trajectory exploration directly in mixed reality, allowing the operator to identify candidate access directions toward the target nodule. In addition, the application provided an estimate of the insertion depth and a first prototype for camera-based needle tracking, with feedback on lateral and angular deviation between the tracked instrument and the planned trajectory.

The phantom-based test confirmed the technical feasibility of the proposed registration workflow. Although the automatic QR-code-based alignment showed an initial vertical offset, the use of a calibration volume allowed the model to be manually refined and satisfactorily aligned with the physical setup. This result suggests that the combination of marker-based registration and calibration-volume-based refinement may represent a practical strategy for future phantom and pre-clinical experiments.

The clinical user questionnaire provided additional preliminary feedback. The most positive responses concerned anatomical visualization, spatial understanding, and laser-based trajectory planning. In particular, the mixed-reality model was perceived as useful for understanding the relationship between the target nodule and the surrounding structures, while the virtual laser was considered promising for supporting trajectory planning. The standardized usability questionnaires also indicated that the system was close to an acceptable usability level, while still requiring improvements to reduce cognitive workload and simplify interaction.

At the same time, the evaluation highlighted important limitations. The study was performed on a single simplified phantom and involved a small number of clinical users; therefore, the results cannot be interpreted as clinical validation. Moreover, the current implementation still requires improvements in registration accuracy, fine model manipulation, needle tracking robustness, rendering optimization, and validation in more realistic experimental conditions.

Overall, the results of this thesis demonstrate the preliminary technical feasibility of an integrated Meta Quest 3-based mixed-reality workflow for thoracic biopsy planning and guidance. The proposed system does not yet represent a clinically validated navigation platform, but it provides a modular foundation for future developments. With further improvements in calibration, tracking, respiratory motion compensation, tissue deformation modelling, and clinical usability, mixed reality on accessible video pass-through headsets such as Meta Quest 3 may become a promising tool for supporting patient-specific planning and guidance in CT-guided thoracic interventions.

Bibliography

[1] Lorenzo Saggiante et al. «Computed-Tomography-Guided Lung Biopsy: A Practice-Oriented Document on Techniques and Principles and a Review of the Literature». en. In: *Diagnostics* 14.11 (May 2024), p. 1089. ISSN: 2075-4418. DOI: 10.3390/diagnostics14111089. URL: https://www.mdpi.com/2075-4418/14/11/1089 (cit. on pp. 5, 6, 10, 15).

[2] David M. DiBardino, Lonny B. Yarmus, and Roy W. Semaan. «Transthoracic needle biopsy of the lung». In: *Journal of Thoracic Disease* 7.Suppl 4 (2015). ISSN: 2077-6624. URL: https://jtd.amegroups.org/article/view/6159 (cit. on pp. 5, 9, 10).

[3] Ronald Winokur, Bradley Pua, Brian Sullivan, and David Madoff. «Percutaneous Lung Biopsy: Technique, Efficacy, and Complications». en. In: *Seminars in Interventional Radiology* 30.02 (May 2013), pp. 121–127. ISSN: 0739-9529, 1098-8963. DOI: 10.1055/s-0033-1342952. URL: http://www.thieme-connect.de/DOI/DOI?10.1055/s-0033-1342952 (cit. on pp. 5, 6, 10).

[4] Ludmila Guralnik, Radu Rozenberg, Alex Frenkel, Ora Israel, and Zohar Keidar. «Metabolic PET/CT-Guided Lung Lesion Biopsies: Impact on Diagnostic Accuracy and Rate of Sampling Error». en. In: *Journal of Nuclear Medicine* 56.4 (Apr. 2015), pp. 518–522. ISSN: 0161-5505, 2159-662X. DOI: 10.2967/jnumed.113.131466. URL: http://jnm.snmjournals.org/lookup/doi/10.2967/jnumed.113.131466 (visited on 06/15/2026) (cit. on p. 6).

[5] Kazuhiko Nakamura, Kensuke Matsumoto, Chie Inoue, Eiji Matsusue, and Shinya Fujii. «Computed Tomography-guided Lung Biopsy: A Review of Techniques for Reducing the Incidence of Complications». en. In: *Interventional Radiology* 6.3 (Nov. 2021), pp. 83–92. ISSN: 2432-0935. DOI: 10.22575/interventionalradiology.2021-0012. URL: https://www.jstage.jst.go.jp/article/interventionalradiology/6/3/6_2021-0012/_article (visited on 06/06/2026) (cit. on pp. 6, 9, 10).

[6] Jared Thomas Weinand, Lourens Du Pisanie, Smith Ngeve, Clayton Commander, and Hyeon Yu. «Pneumothorax after computed tomography-guided lung biopsy: Utility of immediate post-procedure computed tomography and one-hour delayed chest radiography». en. In: *PLOS ONE* 18.4 (Apr. 2023). Ed. by Muhammad Imran, e0284145. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0284145. URL: https://dx.plos.org/10.1371/journal.pone.0284145 (visited on 06/14/2026) (cit. on pp. 6, 9, 10).

[7] Evangelia Siomou, Dimitrios K. Filippiadis, Efstathios P. Efstathopoulos, Ioannis Antonakos, and George S. Panayiotakis. «Patient Dose Estimation in Computed Tomography-Guided Biopsy Procedures». en. In: *Journal of Imaging* 9.12 (Nov. 2023), p. 267. ISSN: 2313-433X. DOI: 10.3390/jimaging9120267. URL: https://www.mdpi.com/2313-433X/9/12/267 (visited on 06/13/2026) (cit. on p. 8).

[8] Kimberly G. Kallianos, Brett M. Elicker, Travis S. Henry, Karen G. Ordovas, Janet Nguyen, and David M. Naeger. «Instituting a Low-dose CT-guided Lung Biopsy Protocol». en. In: *Academic Radiology* 23.9 (Sept. 2016), pp. 1130–1136. ISSN: 10766332. DOI: 10.1016/j.acra.2016.05.013. URL: https://linkinghub.elsevier.com/retrieve/pii/S1076633216300927 (visited on 06/18/2026) (cit. on p. 9).

[9] Takumi Kuriyama, Katsuhiro Masago, Yuki Okada, and Nobuyuki Katakami. «Computed tomography-guided lung biopsy: Association between biopsy needle angle and pneumothorax development». en. In: *Molecular and Clinical Oncology* (Nov. 2017). ISSN: 2049-9450, 2049-9469. DOI: 10.3892/mco.2017.1512. URL: http://www.spandidos-publications.com/10.3892/mco.2017.1512 (visited on 06/19/2026) (cit. on p. 10).

[10] Vania Tacher and Hicham Kobeiter. «State of the Art of Image Guidance in Interventional Radiology». en. In: *Journal of the Belgian Society of Radiology* 102.S1 (Nov. 2018), p. 7. ISSN: 2514-8281. DOI: 10.5334/jbsr.1641. URL: http://jbsr.be/articles/10.5334/jbsr.1641/ (visited on 06/07/2026) (cit. on p. 11).

[11] George C. Kagadis, Konstantinos Katsanos, Dimitris Karnabatidis, George Loudos, George C. Nikiforidis, and William R. Hendee. «Emerging technologies for image guidance and device navigation in interventional radiology». en. In: *Medical Physics* 39.9 (Sept. 2012), pp. 5768–5781. ISSN: 0094-2405, 2473-4209. DOI: 10.1118/1.4747343. URL: https://aapm.onlinelibrary.wiley.com/doi/10.1118/1.4747343 (visited on 06/07/2026) (cit. on p. 11).

[12] Mátyás Rédei et al. «Navigation Systems Significantly Improve the Efficiency and Safety of CT-Guided Interventions». en. In: *Life* 16.3 (Mar. 2026), p. 431.

ISSN: 2075-1729. DOI: 10.3390/life16030431. URL: https://www.mdpi.com/2075-1729/16/3/431 (visited on 06/21/2026) (cit. on pp. 11, 15).

[13] Irene Fuhrmann, Ute Probst, Philipp Wiggermann, and Lukas Beyer. «Navigation Systems for Treatment Planning and Execution of Percutaneous Irreversible Electroporation». en. In: *Technology in Cancer Research & Treatment* 17 (Jan. 2018), p. 1533033818791792. ISSN: 1533-0346, 1533-0338. DOI: 10.1177/1533033818791792. URL: https://journals.sagepub.com/doi/10.1177/1533033818791792 (visited on 06/21/2026) (cit. on pp. 11, 15).

[14] Virender Malik, K Uday Bhanu, and Dhanalakshmi B. «A comparison of robotically assisted needle guidance via the PIGA CT system and conventional CT for percutaneous lung biopsies». en. In: *Journal of Medicine in Scientific Research* 6.4 (Jan. 2023). ISSN: 2537-0928. DOI: 10.59299/2537-0928.1034. URL: https://jmisr.researchcommons.org/home/vol6/iss4/6 (visited on 06/21/2026) (cit. on p. 12).

[15] Shadab Maqsood, Tahir Ahmad Khan, Manzoor Hussain, Mudasir Bhat, Tahleel Altaf Shera, Tariq Ahmad Gojwari, and Feroze A. Shaheen. «Robot-Assisted Lung Biopsy: A Safer Approach to Lung Lesions». en. In: *Journal of Clinical Interventional Radiology ISVIR* 07.03 (Dec. 2023), pp. 166–171. ISSN: 2456-4869. DOI: 10.1055/s-0043-1764371. URL: http://www.thieme-connect.de/DOI/DOI?10.1055/s-0043-1764371 (visited on 06/21/2026) (cit. on p. 12).

[16] Yifan Jing, Jian Jing, Jiayi Liu, Jian Zhang, Yong Jin, and Xuming Bai. «The clinical performance of robotic assisted navigation system versus conventional freehand technique for percutaneous transthoracic needle biopsy». en. In: *Scientific Reports* 15.1 (Feb. 2025), p. 5980. ISSN: 2045-2322. DOI: 10.1038/s41598-025-87987-5. URL: https://www.nature.com/articles/s41598-025-87987-5 (visited on 06/21/2026) (cit. on p. 12).

[17] Pierre Durand, Alexandre Moreau-Gaudry, Anne-Sophie Silvent, Julien Frandon, Emilie Chipon, Maud Médici, and Ivan Bricault. «Computer assisted electromagnetic navigation improves accuracy in computed tomography guided interventions: A prospective randomized clinical trial». en. In: *PLOS ONE* 12.3 (Mar. 2017). Ed. by Iratxe Puebla, e0173751. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0173751. URL: https://dx.plos.org/10.1371/journal.pone.0173751 (visited on 06/21/2026) (cit. on pp. 13, 15).

[18] Morgane Lanouzière, Olivier Varbédian, Olivier Chevallier, Loïc Griviau, Kévin Guillen, Romain Popoff, Serge-Ludwig Aho-Glélé, and Romaric Loffroy. «Computed Tomography-Navigation™ Electromagnetic System Compared to Conventional Computed Tomography Guidance for Percutaneous Lung Biopsy: A Single-Center Experience». en. In: *Diagnostics* 11.9 (Aug. 2021),

p. 1532. ISSN: 2075-4418. DOI: 10.3390/diagnostics11091532. URL: https://www.mdpi.com/2075-4418/11/9/1532 (visited on 06/21/2026) (cit. on pp. 13, 14).

[19] Sara Giannatiempo, Giorgio Carpino, Tommasangelo Petitti, Bruno B. Zobel, Rosario F. Grasso, and Eugenio Guglielmelli. «Efficacy and economic impact evaluation of a navigation system for assisted lung biopsy». en. In: *Healthcare Technology Letters* 5.2 (Apr. 2018), pp. 49–53. ISSN: 2053-3713, 2053-3713. DOI: 10.1049/htl.2017.0015. URL: https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/htl.2017.0015 (visited on 06/22/2026) (cit. on p. 14).

[20] Eliodoro Faiella et al. «Augmented reality 3D navigation system for percutaneous CT-guided pulmonary ground-glass opacity biopsies: a comparison with the standard CT-guided technique». en. In: *Journal of Thoracic Disease* 14.2 (Feb. 2022), pp. 247–256. ISSN: 20721439, 20776624. DOI: 10.21037/jtd-21-1285. URL: https://jtd.amegroups.com/article/view/60677/html (visited on 06/13/2026) (cit. on p. 14).

[21] Michael Evans, Saakhi Kang, Abubakr Bajaber, Kyle Gordon, and Charles Martin. «Augmented Reality for Surgical Navigation: A Review of Advanced Needle Guidance Systems for Percutaneous Tumor Ablation». en. In: *Radiology: Imaging Cancer* 7.1 (Jan. 2025), e230154. ISSN: 2638-616X. DOI: 10.1148/rycan.230154. URL: http://pubs.rsna.org/doi/10.1148/rycan.230154 (visited on 06/10/2026) (cit. on pp. 15, 16, 18, 22).

[22] Amaia Iribar-Zabala, Joseba Ruiz-Olalla-Del-Fresno, Inés Rubio-Pérez, Rafael Moreta-Martínez, Andoni Beristain-Iraola, Javier Pascau, and Mónica García-Sevilla. «A Comparative Study of MR-Guided Needle Insertion for Surgical Procedures: Insights From HoloLens 2, Magic Leap 2 and Apple Vision Pro». en. In: *Healthcare Technology Letters* 13.1 (Jan. 2026), e70037. ISSN: 2053-3713, 2053-3713. DOI: 10.1049/htl2.70037. URL: https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/htl2.70037 (visited on 06/22/2026) (cit. on pp. 15, 16, 22).

[23] Ahmed Elsakka, Brian J. Park, Brett Marinelli, Nathaniel C. Swinburne, and Javin Schefflein. «Virtual and Augmented Reality in Interventional Radiology: Current Applications, Challenges, and Future Directions». en. In: *Techniques in Vascular and Interventional Radiology* 26.3 (Sept. 2023), p. 100919. ISSN: 10892516. DOI: 10.1016/j.tvir.2023.100919. URL: https://linkinghub.elsevier.com/retrieve/pii/S1089251623000355 (visited on 06/22/2026) (cit. on pp. 16, 18, 22).

[24] Marco Casini. «Extended Reality for Smart Building Operation and Maintenance: A Review». en. In: *Energies* 15.10 (May 2022), p. 3785. ISSN: 1996-1073. DOI: 10.3390/en15103785. URL: https://www.mdpi.com/1996-1073/15/10/3785 (visited on 06/24/2026) (cit. on p. 16).

[25] Rafal Nowak, Maja Nowak, Marek Rękas, and Mohammad Javed Ali. «Intraoperative use of mixed reality (MR) in humans across surgical disciplines: a major review». en. In: *3D Printing in Medicine* 11.1 (Dec. 2025), p. 59. ISSN: 2365-6271. DOI: 10.1186/s41205-025-00305-7. URL: https://link.springer.com/10.1186/s41205-025-00305-7 (visited on 06/24/2026) (cit. on p. 16).

[26] Jannick P. Rolland and Henry Fuchs. «Optical Versus Video See-Through Head-Mounted Displays in Medical Visualization». en. In: *Presence: Teleoperators and Virtual Environments* 9.3 (June 2000), pp. 287–309. ISSN: 1054-7460. DOI: 10.1162/105474600566808. URL: https://direct.mit.edu/pvar/article/9/3/287-309/18346 (visited on 06/24/2026) (cit. on pp. 16, 17).

[27] Jan Arensmeyer, Benedetta Bedetti, Philipp Schnorr, Jens Buermann, Donatas Zalepugas, Joachim Schmidt, and Philipp Feodorovici. «A System for Mixed-Reality Holographic Overlays of Real-Time Rendered 3D-Reconstructed Imaging Using a Video Pass-through Head-Mounted Display—A Pathway to Future Navigation in Chest Wall Surgery». en. In: *Journal of Clinical Medicine* 13.7 (Apr. 2024), p. 2080. ISSN: 2077-0383. DOI: 10.3390/jcm13072080. URL: https://www.mdpi.com/2077-0383/13/7/2080 (visited on 06/22/2026) (cit. on pp. 16, 18, 21).

[28] Jessica De Souza and Robert Tartz. «Visual perception and user satisfaction in video see-through head-mounted displays: a mixed-methods evaluation». en. In: *Frontiers in Virtual Reality* 5 (June 2024), p. 1368721. ISSN: 2673-4192. DOI: 10.3389/frvir.2024.1368721. URL: https://www.frontiersin.org/articles/10.3389/frvir.2024.1368721/full (visited on 06/24/2026) (cit. on p. 17).

[29] Jakob Wasserthal et al. «TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images». en. In: *Radiology: Artificial Intelligence* 5.5 (Sept. 2023), e230024. ISSN: 2638-6100. DOI: 10.1148/ryai.230024. URL: http://pubs.rsna.org/doi/10.1148/ryai.230024 (visited on 11/17/2025) (cit. on pp. 18, 26, 27).

[30] Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. «nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation». In: *Nature Methods* 18.2 (Feb. 2021), pp. 203–211. ISSN: 1548-7105. DOI: 10.1038/s41592-020-01008-z. URL: https://doi.org/10.1038/s41592-020-01008-z (cit. on p. 18).

[31] Lijia Zhi, Wujun Jiang, Shaomin Zhang, and Tao Zhou. «Deep neural network pulmonary nodule segmentation methods for CT images: Literature review and experimental comparisons». en. In: *Computers in Biology and Medicine* 164 (Sept. 2023), p. 107321. ISSN: 00104825. DOI: 10.1016/j.compbiomed.2023.107321. URL: https://linkinghub.elsevier.com/retrieve/pii/S0010482523007862 (visited on 06/23/2026) (cit. on p. 18).

[32] Stephanie L. Perkins, Brooke Krajancich, Chi-Fu Jeffrey Yang, Brian A. Hargreaves, Bruce L. Daniel, and Mark F. Berry. «A Patient-Specific Mixed-Reality Visualization Tool for Thoracic Surgical Planning». en. In: *The Annals of Thoracic Surgery* 110.1 (July 2020), pp. 290–295. ISSN: 00034975. DOI: 10.1016/j.athoracsur.2020.01.060. URL: https://linkinghub.elsevier.com/retrieve/pii/S0003497520302782 (visited on 11/14/2025) (cit. on pp. 18, 21).

[33] Ryszard Wierzbicki, Maria Pawłowicz, Józefa Job, Robert Balawender, Wojciech Kostarczyk, Maciej Stanuch, Krzysztof Janc, and Andrzej Skalski. «3D mixed-reality visualization of medical imaging data as a supporting tool for innovative, minimally invasive surgery for gastrointestinal tumors and systemic treatment as a new path in personalized treatment of advanced cancer diseases». en. In: *Journal of Cancer Research and Clinical Oncology* 148.1 (Jan. 2022), pp. 237–243. ISSN: 0171-5216, 1432-1335. DOI: 10.1007/s00432-021-03680-w. URL: https://link.springer.com/10.1007/s00432-021-03680-w (visited on 11/14/2025) (cit. on p. 18).

[34] Tabea Borde et al. «Smart goggles augmented reality CT–US fusion compared to conventional fusion navigation for percutaneous needle insertion». en. In: *International Journal of Computer Assisted Radiology and Surgery* 20.1 (May 2024), pp. 107–115. ISSN: 1861-6429. DOI: 10.1007/s11548-024-03148-5. URL: https://link.springer.com/10.1007/s11548-024-03148-5 (visited on 06/22/2026) (cit. on pp. 21, 22).

[35] Zuodong Song et al. «Augmented reality guided versus computed tomography guided percutaneous lung nodule localization: a noninferiority randomized clinical trial». en. In: *International Journal of Surgery* 111.4 (Apr. 2025), pp. 2933–2941. ISSN: 1743-9159. DOI: 10.1097/JS9.0000000000002308. URL: https://journals.lww.com/10.1097/JS9.0000000000002308 (visited on 06/10/2026) (cit. on pp. 21, 22).

[36] Philip M. Edgcumbe, Tony Jiang, Joshua F. Ho, Michael L. Martin, Michael J. Stein, and Kathryn V. Isaac. «Intraoperative Surgical Guidance for DIEP Flap Harvest Using Augmented Reality». en. In: *Plastic & Reconstructive Surgery* 155.6 (June 2025), pp. 932–935. ISSN: 0032-1052, 1529-4242. DOI:

10.1097/PRS.0000000000011863. URL: https://journals.lww.com/10.1097/PRS.0000000000011863 (visited on 06/24/2026) (cit. on p. 21).

[37] Alexander S. Misono. «Augmented Reality Visualization and Navigation Operationalized in Biopsy of Indeterminate Splenic Mass». en. In: *Clinical Case Reports* 14.1 (Jan. 2026), e71850. ISSN: 2050-0904, 2050-0904. DOI: 10.1002/ccr3.71850. URL: https://onlinelibrary.wiley.com/doi/10.1002/ccr3.71850 (visited on 06/22/2026) (cit. on p. 21).

[38] Ming Li et al. «HoloLens augmented reality system for transperineal free-hand prostate procedures». en. In: *Journal of Medical Imaging* 10.02 (Mar. 2023). ISSN: 2329-4302. DOI: 10.1117/1.JMI.10.2.025001. URL: https://www.spiedigitallibrary.org/journals/journal-of-medical-imaging/volume-10/issue-02/025001/HoloLens-augmented-reality-system-for-transperineal-free-hand-prostate-procedures/10.1117/1.JMI.10.2.025001.full (visited on 06/22/2026) (cit. on p. 22).

[39] Microsoft. *HoloLens 2 release notes*. 2026. URL: https://learn.microsoft.com/en-us/hololens/hololens-release-notes (cit. on p. 23).

[40] Samuel G. Armato et al. «The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A Completed Reference Database of Lung Nodules on CT Scans». en. In: *Medical Physics* 38.2 (Feb. 2011), pp. 915–931. ISSN: 0094-2405, 2473-4209. DOI: 10.1118/1.3528204. URL: https://aapm.onlinelibrary.wiley.com/doi/10.1118/1.3528204 (visited on 11/14/2025) (cit. on p. 25).

[41] 3D Slicer Community. *3D Slicer Official Website*. 2026. URL: https://www.slicer.org/ (cit. on p. 26).

[42] 3D Slicer Community. *3D Slicer Extensions Index*. 2026. URL: https://slicer.readthedocs.io/en/latest/developer_guide/extensions.html (cit. on p. 26).

[43] 3D Slicer Community. *3D Slicer PAA Module*. 2014. URL: https://www.slicer.org/wiki/Documentation/Nightly/Extensions/PercutaneousApproachAnalysis (cit. on p. 26).

[44] 3D Slicer Community. *3D Slicer OpenIGTLinkIF Module*. 2016. URL: https://github.com/openigtlink/SlicerOpenIGTLink (cit. on p. 26).

[45] Jakob Wasserthal and TotalSegmentator contributors. *TotalSegmentator: Tool for robust segmentation of important anatomical structures in CT and MR images*. https://github.com/wasserth/TotalSegmentator. GitHub repository. Accessed: 2026-06-29. 2026 (cit. on p. 27).

[46] Unity Technologies. *Unity Engine: 2D & 3D Development Platform*. https://unity.com/products/unity-engine. Official Unity website. Accessed: 2026-06-28. 2026 (cit. on p. 28).

[47] Meta. *Set up Unity for VR development*. https://developers.meta.com/horizon/documentation/unity/unity-project-setup/. Meta Horizon OS Developers documentation. Accessed: 2026-06-28. 2026 (cit. on p. 28).

[48] Meta. *Meta XR Core SDK (UPM)*. https://developers.meta.com/horizon/downloads/package/meta-xr-core-sdk/. Meta Horizon OS Developers package documentation. Accessed: 2026-06-28. 2026 (cit. on p. 28).

[49] Unity Technologies and glTFast contributors. *Unity glTFast Documentation*. https://github.com/Unity-Technologies/com.unity.cloud.gltfast/blob/main/Packages/com.unity.cloud.gltfast/Documentation~/index.md. GitHub repository documentation. Accessed: 2026-06-28. 2026 (cit. on p. 28).

[50] Meta. *Track QR codes in MR Utility Kit*. https://developers.meta.com/horizon/documentation/unity/unity-mr-utility-kit-qrcode-detection/. Meta Horizon OS Developers documentation. Accessed: 2026-06-28. 2026 (cit. on p. 28).

[51] Meta. *Get Started with Meta XR Simulator*. https://developers.meta.com/horizon/documentation/unity/xrsim-getting-started/. Meta Horizon OS Developers documentation. Accessed: 2026-06-28. 2025 (cit. on p. 28).

[52] Meta. *Meta Quest 3: Next-Gen Virtual Reality Headset*. https://www.meta.com/quest/quest-3/. Official Meta product page. Accessed: 2026-06-28. 2026 (cit. on pp. 30–32).

[53] Meta. *Use Link for App Development*. https://developers.meta.com/horizon/documentation/unity/unity-link/. Meta Horizon OS Developers documentation. Accessed: 2026-06-28. 2026 (cit. on pp. 30, 32).

[54] Meta. *Passthrough Camera API Overview*. https://developers.meta.com/horizon/documentation/spatial-sdk/spatial-sdk-pca-overview/. Meta Horizon OS Developers documentation. Accessed: 2026-06-28. 2026 (cit. on pp. 32, 56).

[55] Meta. *Depth API Overview*. https://developers.meta.com/horizon/documentation/unity/unity-depthapi-overview/. Meta Horizon OS Developers documentation. Accessed: 2026-06-28. 2026 (cit. on p. 32).

[56] Junichi Tokuda et al. «OpenIGTLink: an open network protocol for image-guided therapy environment». en. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 5.4 (Dec. 2009), pp. 423–434. ISSN: 1478-5951, 1478-596X. DOI: 10.1002/rcs.274. URL: https://onlinelibrary.wiley.com/doi/10.1002/rcs.274 (visited on 06/29/2026) (cit. on pp. 31, 52).

[57] Slicer Wiki Contributors. *OpenIGTLinkIF Module*. https://www.slicer.org/wiki/Modules:OpenIGTLinkIF. 3D Slicer documentation. Accessed: 2026-06-29. 2019 (cit. on pp. 33, 45).

[58] Alicia Pose-Díez-de-la-Lastra, Tamas Ungi, David Morton, Gabor Fichtinger, and Javier Pascau. «Real-time integration between Microsoft HoloLens 2 and 3D Slicer with demonstration in pedicle screw placement planning». en. In: *International Journal of Computer Assisted Radiology and Surgery* 18.11 (June 2023), pp. 2023–2032. ISSN: 1861-6429. DOI: 10.1007/s11548-023-02977-0. URL: https://link.springer.com/10.1007/s11548-023-02977-0 (visited on 11/14/2025) (cit. on p. 33).

[59] Slicer Wiki Contributors. *Volume Reslice Driver Module*. https://www.slicer.org/wiki/Documentation/Nightly/Modules/VolumeResliceDriver. 3D Slicer documentation. Accessed: 2026-06-29. 2019 (cit. on pp. 33, 45, 53).

[60] João Cartucho, Chiyu Wang, Baoru Huang, Dan S. Elson, Ara Darzi, and Stamatia Giannarou. «An enhanced marker pattern that achieves improved accuracy in surgical tool tracking». en. In: *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization* 10.4 (July 2022), pp. 400–408. ISSN: 2168-1163, 2168-1171. DOI: 10.1080/21681163.2021.1997647. URL: https://www.tandfonline.com/doi/full/10.1080/21681163.2021.1997647 (visited on 05/05/2026) (cit. on pp. 33, 34, 56).

[61] Pallets Projects. *Flask Documentation*. https://flask.palletsprojects.com/. Official Flask documentation. Accessed: 2026-06-29. 2026 (cit. on p. 35).

[62] OpenCV. *OpenCV: Open Source Computer Vision Library*. https://opencv.org/about/. Official OpenCV website. Accessed: 2026-06-29. 2026 (cit. on p. 35).

[63] Open3D. *open3d.t.geometry.RaycastingScene*. https://www.open3d.org/docs/latest/python_api/open3d.t.geometry.RaycastingScene.html. Official Open3D documentation. Accessed: 2026-06-29. 2018 (cit. on p. 43).

[64] OpenCV. *Perspective-n-Point (PnP) Pose Computation*. https://docs.opencv.org/4.x/d5/d1f/calib3d_solvePnP.html. Official OpenCV documentation. Accessed: 2026-06-29. 2026 (cit. on p. 56).

[65] John Brooke. «SUS: A Quick and Dirty Usability Scale». In: *Usability Evaluation in Industry*. Ed. by Patrick W. Jordan, Bruce Thomas, Ian Lyall McClelland, and Bernard Weerdmeester. Taylor and Francis, 1996, pp. 189–194 (cit. on p. 60).

[66] Sandra G. Hart and Lowell E. Staveland. «Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research». In: *Advances in Psychology*. Vol. 52. Elsevier, 1988, pp. 139–183 (cit. on p. 60).