FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Knee Arthroplasty Navigation System Based on Spatial Augmented Reality

Afonso Pais Pereira



Mestrado em Bioengenharia

Supervisor: Prof. Armando Jorge Miranda de Sousa (FEUP)

Co-Supervisor: Prof. Joaquim Gabriel Magalhães Mendes (FEUP)

Collaborator: Dr. Adélio Justino Machado Vilaça (ULSSA)

Collaborator: Catarina Oliveira Lopes (FEUP)

September 06, 2026

Knee Arthroplasty Navigation System Based on Spatial Augmented Reality

Afonso Pais Pereira

Mestrado em Bioengenharia

September 06, 2026

Abstract

Total knee arthroplasty (TKA) is among the most frequently performed orthopaedic procedures worldwide, yet a notable proportion of patients experience persistent post-operative dissatisfaction, often stemming from suboptimal implant alignment. While robotic-assisted surgical systems significantly enhance resection accuracy and component positioning, their widespread adoption remains constrained by prohibitive capital costs, substantial operating room footprint, and added workflow complexity. Spatial augmented reality (SAR) presents a compelling alternative capable of superimposing navigational cues directly onto the patient anatomy without requiring bulky robotic arms or cumbersome head-mounted displays.

This work presents the design, implementation, and benchtop experimental validation of a markerless SAR navigation framework for distal femoral resection in TKA. A dedicated surgical planning application was developed to define the patient-specific resection parameters and determine the target placement and orientation of the distal cutting guide based on pre-operative Computed Tomography (CT) models. The system couples an Ensenso N36 stereo camera with a digital projector, calibrated using structured light patterns to establish direct correspondence between projector pixel space and three-dimensional (3D) physical coordinates. Intraoperatively, dense point clouds acquired from the exposed femoral surface are registered to patient-specific models segmented from pre-operative CT scans using coarse-to-fine point cloud registration algorithms. Through this calibrated spatial transformation, the system projects the target position and orientation of the distal femoral cutting guide into the surgical field in real time, designed to visually align with the upper surface of the physical guide once it achieves its planned surgical pose.

The framework was experimentally validated on 3D-printed femoral phantoms across calibration, surface registration, and physical guide positioning. Projector–camera calibration verified accurate closed-loop optical overlay, while the multi-stage registration achieved sub-millimeter inlier root-mean-square error (RMSE) between 0.34 mm and 0.56 mm across full and partially occluded bone surfaces. In physical positioning trials, the system yielded an overall relative positioning mean absolute error (MAE) of 3.88 ± 3.46 mm and an angular MAE of $4.91 \pm 3.85^{\circ}$. While further mechanical and algorithmic refinements are needed to meet strict surgical tolerances, these benchtop results highlight the viability of markerless SAR as a cost-effective, space-efficient alternative to conventional instrumentation and robotic systems.

Keywords: Point Cloud Registration, Projector-Camera Calibration, Spatial Augmented Reality, Surgical Navigation, Surgical Planning, Total Knee Arthroplasty

Resumo

A artroplastia total do joelho (ATJ) é um dos procedimentos ortopédicos mais frequentemente realizados a nível mundial, contudo uma proporção notável de pacientes experiencia insatisfação pós-operatória persistente, frequentemente decorrente de um alinhamento subótimo do implante. Embora os sistemas cirúrgicos assistidos por robótica aumentem significativamente a precisão da resseção e o posicionamento dos componentes, a sua ampla adoção permanece limitada pelos custos de aquisição proibitivos, pela espaço ocupado no bloco operatório e pelo aumento da complexidade do fluxo de trabalho. A realidade aumentada (RA) surge como uma alternativa apelativa, capaz de sobrepor indicações visuais de navegação diretamente na anatomia do paciente, sem a necessidade de braços robóticos volumosos ou equipammentos pessoais desconfortáveis.

Este trabalho apresenta o desenvolvimento, implementação e validação experimental de um sistema de navegação em RA para a resseção femoral distal em ATJ. Foi desenvolvida uma aplicação dedicada de planeamento cirúrgico para definir os parâmetros de resseção específicos do paciente e determinar o posicionamento e orientação planeados para a guia de corte distal, com base em modelos pré-operatórios de Tomografia Computorizada (TC). O sistema utiliza uma sensor de profundidade Ensenso N36 com um projetor digital, calibrados através de padrões de *structured light* de modo a estabelecer uma correspondência direta entre os píxeis do projetor e as coordenadas físicas tridimensionais (3D). No intraoperatório, nuvens de pontos adquiridas a partir da superfície femoral exposta são registadas nos modelos específicos do paciente, segmentados a partir das imagens de TC pré-operatórias, através de algoritmos de registo de nuvens de pontos com abordagem *coarse-to-fine*. Através desta transformação espacial calibrada, o sistema projeta a posição e orientação alvo da guia de corte femoral distal no campo cirúrgico em tempo real, concebida para se alinhar visualmente com a superfície superior da guia física assim que esta atinge a posição cirúrgica planeada.

A metodologia foi validada experimentalmente em modelos anatómicos femorais impressos em 3D, englobando a calibração, o registo de superfícies e o posicionamento do guia físico. A calibração projetor-câmara confirmou uma sobreposição ótica precisa em regime (*closed-loop*), enquanto o registo multi-etapa alcançou um *root-mean-square error* (RMSE) submilimétrico de *inliers* entre 0,34 mm e 0,56 mm, tanto em superfícies ósseas completas como parcialmente ocluídas. Nos ensaios de posicionamento físico, o sistema obteve um *mean absolute error* (MAE) de posicionamento relativo global de 3,88 $\pm$ 3,46 mm e um MAE angular de 4,91 $\pm$ 3,85°. Embora sejam necessários refinamentos mecânicos e algorítmicos adicionais para cumprir consistentemente as estritas tolerâncias cirúrgicas, estes resultados destacam a viabilidade da RA como uma alternativa economicamente acessível e eficiente ao nível de espaço em relação à instrumentação convencional e aos sistemas robóticos.

Keywords:Artroplastia Total do Joelho , Calibração Projetor-Câmara, Navegação Cirúrgica, Planeamento Cirúrgico, Realidade Aumentada, Registo de Nuvens de Pontos

Acknowledgements

I would like to express my sincere gratitude to my supervisors, Prof. Armando Sousa and Prof. Joaquim Mendes, for their continuous guidance, support, and trust throughout this work.

A special thank you goes to Catarina Lopes for her close collaboration and invaluable contributions. I also thank Dr. Adélio Vilaça for his collaboration and helpful input.I am also grateful to the laboratory staff and my colleagues for always being ready to help and for creating such a great environment in the lab.

Finally, a special thank you to my friends and, above all, to my family for their unconditional love, patience, and encouragement. Without their presence every step of the way, this would not have been possible.

Afonso Pereira

"Courage isn't having the strength to go on
it is going on when you don't have strength."

Napoleon Bonaparte

Contents

1 Introduction **1**
- 1.1 Context and Motivation . . . 1
- 1.2 Main Objectives . . . 2
- 1.3 Sustainable Development Goals . . . 2
- 1.4 Scientific Contributions . . . 3
- 1.5 Document Structure . . . 3

2 Anatomy of the Knee **5**
- 2.1 Joint Classification . . . 5
- 2.2 Overview of the Knee . . . 5
  - 2.2.1 Bony Structure . . . 6
  - 2.2.2 Muscle Structure . . . 7
  - 2.2.3 Menisci . . . 7
  - 2.2.4 Ligaments . . . 8
- 2.3 Knee Pathologies . . . 10
- 2.4 Total Knee Arthroplasty Surgery . . . 11
  - 2.4.1 Pre-operative Evaluation and Risk Stratification . . . 11
  - 2.4.2 Surgical Exposures and Approaches . . . 13

3 State of Art **19**
- 3.1 Point Cloud Registration . . . 19
  - 3.1.1 Rigid Point Cloud Registration Algorithms . . . 19
  - 3.1.2 Deep Learning-Based Registration Paradigms . . . 21
  - 3.1.3 Point Cloud Registration in Total Knee Arthroplasty . . . 22
- 3.2 Projection Mapping . . . 23
  - 3.2.1 Static Projector-Camera Calibration Fundamentals . . . 24
  - 3.2.2 Dynamic Projection Mapping . . . 25
  - 3.2.3 Rigid Projections . . . 25
  - 3.2.4 Non-Rigid Projections . . . 26

4 Methodology **31**
- 4.1 Development Workflow . . . 31
- 4.2 System Workflow . . . 32
  - 4.2.1 Equipment, Software and Tools . . . 33
- 4.3 Surgical Planning Application . . . 33
- 4.4 Surgical Planning Pipeline . . . 34
  - 4.4.1 Femur 3D model . . . 34
  - 4.4.2 Anatomical Landmark Definition and Axis Alignment . . . 35

4.5 Point Cloud Registration . . . 42
4.5.1 Data Acquisition System . . . 42
4.5.2 Coarse Alignment . . . 43
4.5.3 Fine Registration . . . 45
4.5.4 Real-Time Tracking . . . 46
4.5.5 Validation . . . 47
4.6 Projector-Camera Calibration . . . 50
4.6.1 Intrinsic Parameters and the Pinhole Model . . . 50
4.6.2 Planar Calibration and Zhang's Methodology . . . 50
4.6.3 Lens Distortion and Image Rectification . . . 51
4.6.4 Ensenso Depth Camera Intrinsic Parameters . . . 51
4.6.5 Projector Calibration . . . 52
4.7 Spatial Augmented Navigation System Implementation . . . 55
4.7.1 Hardware Setup and Configuration . . . 55
4.7.2 Surgical Guide Support Configuration . . . 57
4.7.3 Intraoperative Alignment, Tracking, and Projection Pipeline . . . 57
4.7.4 Pipeline Validation Methodology . . . 61

5 Results and Discussion **65**
5.1 Camera Error Analysis . . . 65
5.2 Registration Pipeline Validation . . . 66
5.2.1 Phase 1: Controlled System Validation . . . 66
5.2.2 Phase 2: Clinical Applicability and Inter-Operator Reproducibility . . . 67
5.2.3 Phase 3: Direct Surgical Target Validation on Femur Morphology . . . 68
5.2.4 Current Limitations and Translational Challenges . . . 70
5.3 Projector–Camera Calibration and Validation . . . 71
5.4 SAR Navigation System . . . 73

6 Conclusions and Future Work **79**

A Results Tables **81**
A.1 Raw Benchtop Coordinate Measurements . . . 81

References **83**

List of Figures

2.1 Anatomy of the Knee Joint, adapted from [1] . . . 6
2.2 Musculature surrounding the knee joint [2] . . . 7
2.3 Anatomical overview of the medial and lateral menisci of the knee, adapted from [3] 8
2.4 Anterior (reflected patella) and posterior anatomical views of the right knee joint, detailing primary stabilizing ligaments, menisci, and articular bone landmarks[4] 9
2.5 Femoral resection angle defined between the anatomical and mechanical axes of the femur [5] . . . 12
2.6 Alternative surgical approaches to the knee [5] . . . 13
2.7 Initial steps of the femoral preparation [5] . . . 14
2.8 Positioning and securing of the distal femoral cutting guide [5] . . . 15
2.9 Distal femoral resection using a rotating saw [5] . . . 15
2.10 Sizing guide attached to the distal femur [5] . . . 16
2.11 Positioning of the 4-in-1 cutting guide and execution of the femoral resections [5] 16
2.12 Tibia cutting guide [5] . . . 17

4.1 End-to-end operational workflow of the navigation system, encompassing pre-operative planning and intraoperative registration and projection . . . 32
4.2 Surgical planning application user interface layout showing the 3D visualization area (red), 2D visualization area (yellow), and side panel (white) . . . 33
4.3 Surgical planning interface and available toolbar options . . . 34
4.4 Definition of the femoral anatomical axis (red) via proximal and distal landmark selection . . . 35
4.5 Simultaneous visualization of the virtual mechanical axis (green) and the anatomical axis (red) . . . 36
4.6 Calculated angular deviation between the mechanical and intramedullary axes, representing the projected distal cut angle . . . 36
4.7 Definition of the cutting plane origin and corresponding side panel interface . . . 36
4.8 Definition of the cutting plane using the furthest point and corresponding side panel interface . . . 37
4.9 Rotational adjustment of the cutting plane and corresponding side panel interface 38
4.10 Inferior and superior 2D views of the planned cutting plane . . . 39
4.11 Workspace layout inversion maximizing the 2D planning window for detailed cross-sectional inspection . . . 39
4.12 Distance measurement functionality available in the 2D planning view . . . 40
4.13 Reference point annotation feature in the 2D visualization environment . . . 40
4.14 Virtual seating simulation of the localized femoral distal cutting block onto the bone mesh surface . . . 41
4.15 Front and back views of the distal femoral cutting guide . . . 42

4.16 NxView calibration wizard interface during the camera validation routine 43
4.17 Pre-operative landmark selection application used to identify anatomical reference points on the femur model 44
4.18 Intraoperative landmark matching interface showing the pre-operative femur model with selected landmarks (left) and the acquired point cloud (right) 45
4.19 Point clouds after coarse alignment using the initial transformation estimate. . . . 45
4.20 Point clouds after the ICP refinement 46
4.21 Graphical user interface used during real-time tracking, displaying the continuously the transformation matrix between the pre-operative model and the live point cloud 47
4.22 Validation rig used in the study 48
4.23 Experimental surface visibility conditions of the 3D-printed femurs used for registration validation 49
4.24 Chess board pattern 52
4.25 Projected Gray-code structured light pattern captured by the left camera 53
4.26 Operating room configuration and equipment layout during conventional TKA [6] 55
4.27 Experimental hardware setup for the SAR system, showing the side-by-side mounting of the projector and Ensenso N36 depth camera above the femoral target . . . 56
4.28 Mechanical support assembly for surgical guide stabilization, consisting of an articulating mechanical arm and a 3D-printed adapter 57
4.29 Boundary point clouds overlaid on the distal femoral cutting block model, showing the outer boundary (P_{outer}) highlighted in red and the inner boundary (P_{inner}) highlighted in green 58
4.30 Window showing the 2D image of the surgical guide borders being projected . . 60
4.31 Comparison of the projection with and without the guide. 60
4.32 Selected guide points in the distal femoral guide 3D model 61
4.33 3D models of the femur and surgical guide in the planned operational position, showing the selected femoral (red) and guide (yellow) points used to define their expected spatial relationship 62

5.1 Qualitative visual comparison of the registration results between the fully exposed femur (a) and the clinically emulated partially masked femur (b) for model set 1 . 70
5.2 Qualitative visual comparison of the registration results between the fully exposed femur (a) and the clinically emulated partially masked femur (b) for model set 2 . 70
5.3 Qualitative visual comparison of the registration results between the fully exposed femur (a) and the clinically emulated partially masked femur (b) for model set 3 . 70
5.4 Closed-loop back-projection validation 72

List of Tables

1.1 Strategic alignment of the SAR navigation framework with the UN Sustainable Development Goals . . . 3

3.1 Summary of local refinement, global alignment, and deep learning-based point cloud registration methods . . . 28
3.2 Summary of static calibration and dynamic projection mapping methods discussed in the literature . . . 29

5.1 Tracking validation results . . . 65
5.2 Quantitative registration metrics before and after point-to-plane ICP Refinement . 66
5.3 Quantitative alignment metrics of the user study evaluating the registration pipeline on the external knee surface ($n = 15$) . . . 67
5.4 Alignment metrics for femur models . . . 68
5.5 Optimized projector intrinsic and distortion parameters. . . . 71
5.6 System-wide summary of end-to-end SAR projection and positioning errors across n=3 femoral specimens. Values are reported as mean ± SD . . . 73
5.7 Absolute guide point 3D residual displacement errors after rigid registration . . . 73
5.8 Surgical navigation verification: comprehensive planned vs. measured deviations for specimen set 1 (Femur 1). . . . 76
5.9 Surgical Navigation Verification: Comprehensive planned vs. measured deviations for specimen set 2 (Femur 2). . . . 76
5.10 Surgical navigation verification: comprehensive planned vs. measured deviations for specimen set 3 (Femur 3). . . . 77

A.1 Femur 1: 3D coordinates of the femur (F) and guide points(G) . . . 81
A.2 Femur 2: 3D coordinates of the femur (F) and guide points(G) . . . 82
A.3 Femur 3: 3D coordinates of the femur (F) and guide points(G) . . . 82

Abbreviations

2D	Two-dimensional
3D	Three-dimensional
ACL	Anterior cruciate ligament
CAD	Computer-aided Design
CT	Computed Tomography
DPM	Dynamic Projection Mapping
FPFH	Fast Point Feature Histograms
FRE	Fiducial Registration Error
ICP	Iterative Closest Point
IR	Infrared
JRMPC	Joint Robust Point Cloud Registration
LCL	Lateral collateral ligament
LiDAR	Light Detection and Ranging
MAE	Mean Absolute Error
MCL	Medial collateral ligament
MRI	Magnetic Resonance Imaging
NDT	Normal Distributions Transform
PCL	Posterior cruciate ligament
RANSAC	Random Sample Consensus
RMSE	Root Mean Square Error
ROS2	Robot Operating System 2
SAR	Spatial Augmented Reality
SDG	Sustainable Development Goals
SME	Signed Mean Error
STL	Stereolithography
TKA	Total Knee Arthroplasty
TRE	Target Registration Error
VMO	Vastus medialis obliquus

Chapter 1

Introduction

Total knee arthroplasty (TKA) is a highly prevalent orthopedic procedure, with approximately 3.6 million surgeries performed globally each year [7], a number projected to surge due to demographic shifts and lifestyle risk factors [8]. Despite the procedure's general success, up to 19% of patients remain dissatisfied with their post-operative outcomes [9], primarily due to suboptimal implant alignment. To improve accuracy, robotic-assisted navigation systems were introduced, however they are expensive and physically intrusive, disrupting standard surgical workflows [10]. To overcome these limitations, this work aims to develop a novel Spatial Augmented Reality (SAR) navigation system that delivers robotic-level precision in a cost-effective, compact format, seamlessly integrating into the operating room without disrupting the surgical process.

1.1 Context and Motivation

TKA, commonly referred to as total knee replacement surgery, is a major orthopedic procedure designed to replace damaged bone and cartilage within the knee joint with artificial biomaterials [11]. The primary clinical objectives of TKA are to alleviate severe, chronic pain and to restore functional mobility, thereby directly improving the patient's overall quality of life [12]. Globally, TKA is one of the most frequently performed surgeries, with millions of procedures executed annually worldwide, a figure projected to grow exponentially over the coming decades [8]. This surge is heavily driven by a rapidly aging global population and rising obesity rates, both of which significantly accelerate mechanical wear on weight-bearing joints [8]. The intervention is typically indicated when conservative treatments fail to manage end-stage joint degeneration caused by underlying pathologies [13]. The primary disease necessitating this surgery is osteoarthritis, a progressive "wear-and-tear" breakdown of articular cartilage, though it is also frequently indicated for rheumatoid arthritis [14], a systemic autoimmune disorder, and post-traumatic arthritis resulting from severe joint injuries [15].

The surgical intervention typically lasts between one and three hours [16], during which the damaged articular surfaces from the distal end of the femur and the proximal end of the tibia are surgically resected [11]. This bone removal is executed to establish two flat surfaces that are perpendicular to the mechanical axis of the lower limb, an orientation that is paramount to restoring the natural alignment of the leg and ensuring even weight distribution across the joint [13]. High-grade biocompatible metal components are then securely fixed to the newly shaped ends of the femur and tibia to replace the resected bone structure [17]. To complete the assembly, a specialized medical-grade plastic spacer is inserted between the metallic components [17]. The plastic surface mimics the low-friction properties of healthy natural cartilage, allowing the metal components to glide smoothly against one another while providing necessary shock absorption and joint stability, which prevents direct metal-on-metal contact that would otherwise cause rapid wear and debris formation [18].

1.2 Main Objectives

To overcome the logistical and financial challenges of robotic-assisted surgery without sacrificing intraoperative precision, this work aims to develop a novel surgical navigation system based on SAR. The proposed system is designed to deliver the high accuracy characteristic of robotics within a cost-effective, compact format that seamlessly integrates into the clinical environment and preserves the surgeon's operative workflow.

To achieve this primary objective, the specific goals of this thesis are structured as follows:

1. **Surgical Planning Application:** Develop a dedicated software platform to visualize patient anatomy and define the target positioning of surgical instruments and resection planes.
2. **Markerless Tracking and Spatial Calibration:** Implement a computer vision framework to calibrate the sensing and projection hardware and align intraoperative surface data with pre-operative anatomical models.
3. **Dynamic Visual Guidance:** Design and integrate a SAR engine that projects planned instrument placement cues directly into the surgical field in real time.

1.3 Sustainable Development Goals

In alignment with the United Nations Educational, Scientific and Cultural Organization report *"Engineering for Sustainable Development: How Engineering Can Make It Happen"*[19], the development of this SAR surgical navigation system leverages engineering innovation to address critical global challenges. By transitioning from high-cost, resource-intensive robotic surgery frameworks to an accessible, software-driven projector-camera architecture, this research directly contributes to the United Nations Sustainable Development Goals (SDG). The structural impact of this project, mapped across specific targets, contributions, and measurable performance indicators, is detailed in Table 1.1.

Table 1.1: Strategic alignment of the SAR navigation framework with the UN Sustainable Development Goals

SDG	Target	Contribution	Performance Indicators & Metrics
Goal 9: Industry, Innovation, and Infrastructure	**9.5** Enhance scientific research and upgrade technological capabilities.	Develops an open, modular spatial augmented reality framework for computer-assisted interventions, advancing biomedical engineering research and technological capacity.	**Indicator:** Computational efficiency and system robustness. • Geometric calibration and projection accuracy. • End-to-end computational latency and execution time. • Reproducibility and modularity across standard computing hardware.
Goal 3: Good Health and Well-being	**3.8** Achieve access to safe, effective, and affordable essential healthcare services.	Lowers technological and economic barriers to computer-assisted orthopedic surgery by providing an accessible alternative to expensive proprietary systems.	**Indicator:** Economic accessibility and surgical guidance precision. • Hardware acquisition and implementation costs relative to commercial alternatives. • Spatial tracking and anatomical surface alignment accuracy. • System usability and workflow feasibility on anatomical models.

1.4 Scientific Contributions

A paper focusing on the registration aspect of this work was accepted for publication and presented at the 13th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI 2026). Additionally, an early version of the research was presented at the *19th Encontro de Investigação Jovem da U.Porto* (IJUP).

1.5 Document Structure

This dissertation is organized into five subsequent chapters following this introduction. Chapter 2 establishes the anatomical framework by detailing the musculoskeletal mechanics of the knee joint, common degradation pathways such as osteoarthritis, and the surgical protocols of TKA. Chapter 3 reviews the technological landscape of surgical guidance, beginning with a analysis of intraoperative surgical registration. It then evaluates contemporary projection mapping methodologies. Chapter 4 presents the technical methodology developed in this work, detailing the engineering of the surgical planning application, the pipelines deployed for rigid Three-dimensional (3D) point cloud alignment, the mathematical models utilized for projector-camera calibration, and the integration of the overall system configuration designed to guide the positioning of surgical cuts. Chapter 5 provides an evaluation of the system, analyzing the performance of the hardware setup, the accuracy of the spatial registration, and the tracking latency under simulated

operating room conditions. Finally, Chapter 6 summarizes the primary contributions of this research, outlines current architectural limitations, and proposes directions for future development.

Chapter 2

Anatomy of the Knee

2.1 Joint Classification

A joint or articulation is the site where two or more adjacent bones, or bone and cartilage connect. Joints can be systematically categorized based on two main criteria: their structure (material connecting the bones or cartilage) and their function (the degree of mobility they allow) [20].

The structural classification of a joint is determined by the manner in which the articulating bone surfaces are connected, specifically whether they are linked by fibrous connective tissue, cartilage, or if they meet within a fluid-filled joint cavity [21]. Structurally, joints are divided into three main categories: synovial joints, that are characterized by a joint cavity filled with a lubricating fluid known as synovial fluid. In this type of joint the articulating bones are not directly connected by solid tissue. Instead, their surfaces meet within this cavity, allowing for a wider range of movement; cartilaginous joints when the bones are directly united by either hyaline cartilage or fibrocartilage; and fibrous joints, when the bones are joined together by dense fibrous connective tissue [20].

The functional classification relates directly to the joint's mobility capacity. Joints are thus functionally classified into three types based on movement: synarthroses, which are immobile or nearly immobile; amphiarthroses, which allow slight movement; and diarthroses, which are freely movable [21]. Structurally, fibrous and cartilaginous joints may correspond to either synarthroses or amphiarthroses, depending on their location and composition. In contrast, all synovial joints are functionally classified as diarthroses, reflecting their capacity for free movement.

2.2 Overview of the Knee

The knee joint is the largest joint in the human body and can be classified as a synovial diarthrosis joint. It has a highly complex structure that allows for the optimal interplay of the bony structures,

ligaments, tendons, muscles, and the joint capsule[22, 23]. The knee joint consists of two articulations enclosed within a single articular capsule: the patellofemoral joint and the femorotibial joint. The patellofemoral joint, formed between the patella and the femur, facilitates the efficient and low-friction transmission of forces generated by the quadriceps femoris muscle. The femorotibial joint, formed between the femur and the tibia, is responsible for supporting the majority of the body's weight[24].

Functionally, the knee is a trochogingylmos joint, meaning it can be classified as a modified gliding hinge joint. While it primarily acts as a hinge, allowing flexion and extension in the sagittal plane, it is also capable of a wider range of motion than a simple hinge, like gliding movement, offering six degrees of freedom of motion including: rotational movement, including flexion-extension, internal–external and varus–valgus; translational movement, in the anterior–posterior and medial–lateral directions; and by compression and distraction of the knee[23, 24].

2.2.1 Bony Structure

The knee joint involves the articulation of three distinct bones: the distal femur, the proximal tibia and the patella, as shown in Figure 2.1.

The height of the femoral condyles exhibits asymmetry, with the medial condyle being both larger and projecting more distally than the lateral femoral condyle. They are separated by the femoral groove anteriorly and the femoral notch distally. A key identifier on the lateral condyle is the terminal sulcus and the groove marking the insertion point of the popliteus muscle[24–26].

The proximal tibia consists of the tibial plateau, which have an approximate 10° posterior slope and are relatively flat. Dividing the plateaus, there is the tibial spine where critical soft tissues attach, including the anterior cruciate ligament (ACL), posterior cruciate ligament (PCL) and the horns of the meniscus[24–26].

The patella is the largest sesamoid bone of the human body, measuring around 5 cm in diameter. It possesses the thickest articular surface in the body to withstand the joint reactive forces that occur between the patella and the femur[24–26].

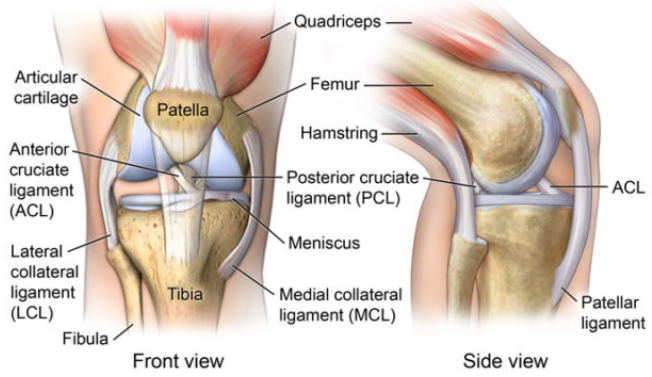


Figure 2.1: Anatomy of the Knee Joint, adapted from [1]

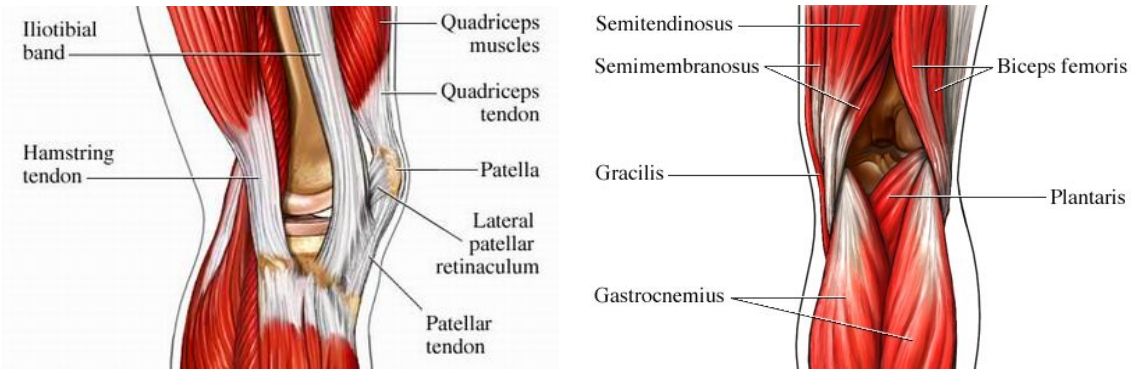
2.2.2 Muscle Structure

The muscles surrounding the knee act as primary motor units driving motion across all six degrees of freedom, while functioning as secondary stabilizers crucial for joint proprioception and neuromuscular control. These muscles are categorized as monoarticular, acting only at the knee, or biarticular, acting at both the hip and knee [24]. The musculature surrounding the knee joint is shown in Figure 2.2.

Anteriorly, knee extension is driven by the quadriceps femoris group, which converges into a trilaminar insertion at the superior patella before forming the 5 cm patellar tendon that attaches to the tibial tubercle [24, 25]. This group consists of the superficial rectus femoris, the middle vastus lateralis and vastus medialis, and the deep vastus intermedius. Within this group, the vastus medialis obliquus (VMO) features fibers oriented at a 55° to 70° angle to control patellar tracking, while the vastus lateralis expands into the lateral retinaculum [25].

Posteriorly, flexion and rotation are governed by the hamstrings and calf musculature. The hamstrings include the lateral-rotating biceps femoris and the medial-rotating semimembranosus and semitendinosus. Secondary flexion is aided by the plantaris and the gastrocnemius heads, while the deep soleus assists by resisting anterior joint translation[24].

Coronal and rotational stability are supported by dedicated medial and lateral groups. Medially, the pes anserinus complex, composed by the sartorius, gracilis, and semitendinosus, works with the semimembranosus to flex and internally rotate the joint [24, 25]. Laterally, the iliotibial band and tensor fasciae latae provide dynamic restraint, assisted by the lateral gastrocnemius and biceps femoris [24, 25]. Finally, the deep popliteus muscle executes internal and external rotation to lock and unlock the knee joint [24, 25].



(a) Lateral view of the musculature of the knee (b) Posterior view of the musculature of the knee

Figure 2.2: Musculature surrounding the knee joint [2]

2.2.3 Menisci

The menisci are two crescent shaped fibrocartilaginous structures, medial and lateral, interposed between the femoral condyles and the tibial plateau, as shown in Figure 2.3. Their primary function is to bear and distribute loads across the knee joint while absorbing shock, and they may

also act as secondary restraints to anterior tibial translation in an ACL deficient knee[24–26]. The menisci are composed primarily of collagen (approximately 75%) and non-collagenous proteins (8–13%). Type I collagen constitutes about 90% of the total collagen content and is predominantly arranged in a circumferential pattern, allowing the menisci to resist hoop stresses generated during compressive loading. Radial collagen fibers provide structural integrity by limiting longitudinal splitting and contributing to overall rigidity[24–26]. Each meniscus covers around two-thirds of the corresponding tibial articular surface, with a thick and convex peripheral border attached to the inside of the capsule joint, while its inner border is thin and concave, and has a free edge. The medial meniscus is more firmly attached to the tibial plateau and joint capsule, whereas the lateral meniscus exhibits greater mobility during knee flexion and rotation[24, 26].

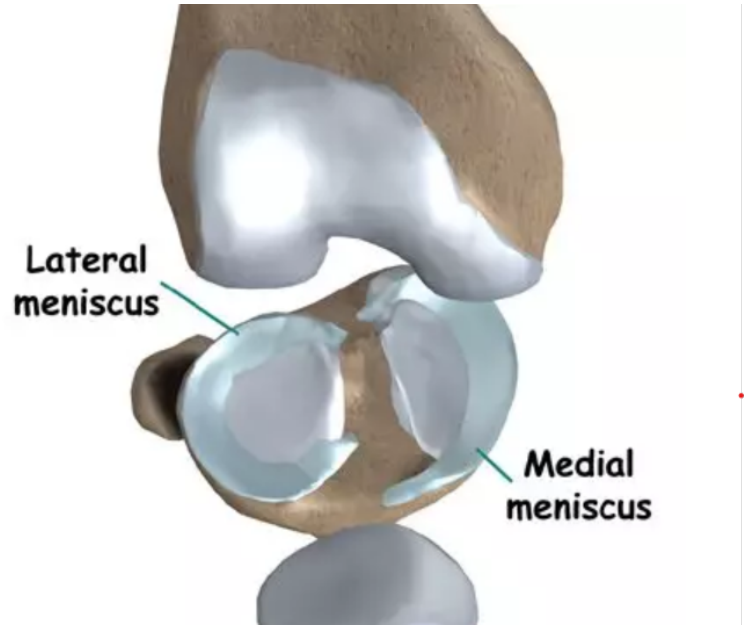


Figure 2.3: Anatomical overview of the medial and lateral menisci of the knee, adapted from [3]

2.2.4 Ligaments

The ligaments are responsible for providing stability, being the primary stabilizers of the knee. They play a critical role in maintaining joint congruency and resisting excessive translation and rotation of the tibia relative to the femur, especially in high-load and extreme-range-of-motion situations[24, 25].

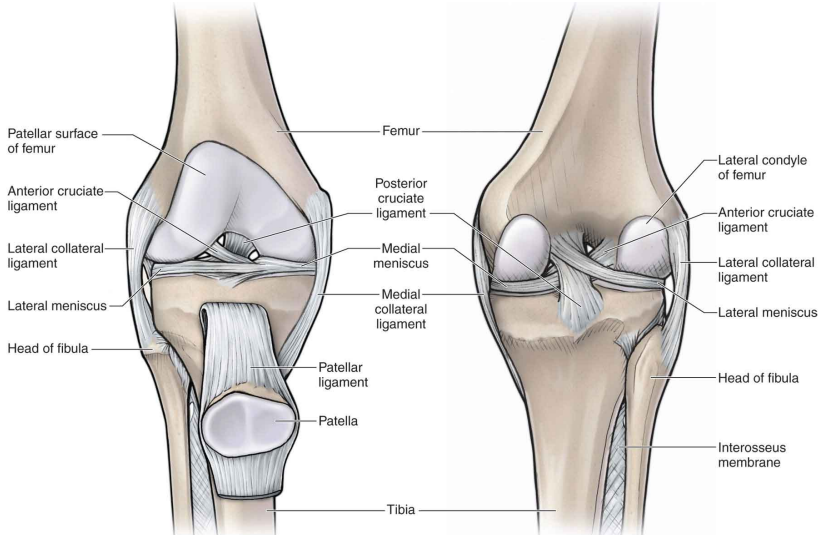


Figure 2.4: Anterior (reflected patella) and posterior anatomical views of the right knee joint, detailing primary stabilizing ligaments, menisci, and articular bone landmarks[4]

2.2.4.1 Cruciate Ligaments

The ACL acts as the primary static restraint against anterior translation of the tibia on the femur and also helps control the rotational stability of the knee. It originates from the anterior intercondylar region of the tibia and inserts on the posteromedial aspect of the lateral femoral condyle. The ACL consists of two functional bundles: the anteromedial bundle becomes taut during knee flexion, while the posterolateral bundle is tight in full extension. The bundles' contribution varies along the knee flexion, with the ACL experiencing minimal strain between 20° and 30° of knee flexion[24–26]. The PCL acts as the primary restraint against the posterior translation of the tibia on the femur, providing approximately 95% of the total resistance to this motion. It is larger and stronger than the ACL, with nearly twice its tensile strength and a greater cross-sectional area. It originates from the lateral surface of the medial femoral condyle and inserts on the posterior intercondylar region of the tibia. The PCL consists of two bundles: the anterolateral bundle, which tightens in flexion, and a smaller posteromedial bundle, which becomes taut in extension. The Humphrey and Wrisberg ligaments may also be present anteriorly and posteriorly to the PCL, contributing as secondary stabilizers against posterior tibial translation[24–26]. Both the ACL and the PCL are intra-articular and extrasynovial, receiving their vascular supply from the middle geniculate artery[24–26].

2.2.4.2 Medial Ligament Complex

The medial collateral ligament (MCL) is the primary stabilizer against valgus stress at the knee [24–26]. The medial side of the knee is organized into three anatomical layers. The superficial

layer contains the sartorius and its fascia, the middle layer includes the superficial MCL and posterior oblique ligament, and the deep layer consists of the joint capsule and the deep MCL, also known as the medial capsular ligament[24–26].

The superficial MCL starts from the medial femoral epicondyle and inserts distally into the medial tibia. It is composed of anterior (parallel) and posterior (oblique) fibers, which display variable tension along the knee motion. The anterior knee fibers have a greater tautness between 45° and 90° of flexion, while the anterior fibers relax with increasing flexion. The superficial MCL resists valgus stress and contributes to control of external rotation of the tibia, while the deep MCL connects the ligamentous complex to the medial meniscus and provides a secondary role in valgus stability[24–26].

2.2.4.3 Lateral Ligament Complex and Posterolateral Stabilizers

The lateral side of the knee is also organized into three layers. The superficial layer includes the iliotibial tract and biceps femoris tendon, while the intermediate layer contains the patellar retinaculum, and the deep layer includes the lateral collateral ligament (LCL), popliteus tendon, popliteofibular ligament, arcuate ligament, fabellofibular ligament, and joint capsule[24–26].

The LCL is a strong, cord-like structure extending from the lateral femoral condyle to the fibular head. It serves as the primary static restraint to varus stress and also resists external rotation of the knee [24–26]. Unlike the MCL, the LCL is extracapsular and does not attach to the lateral meniscus.

The popliteus muscle–tendon unit plays a key role in posterolateral stability by actively internally rotating the tibia to unlock the knee from full extension and passively resisting excessive external tibial rotation. Together with the popliteofibular ligament, arcuate ligament, and fabellofibular ligament, these structures stabilize the posterolateral corner against varus stress, posterior tibial translation, and rotational instability, particularly in conjunction with the PCL [25, 26].

2.3 Knee Pathologies

TKA has been performed for more than 40 years, with it being one of the most common and cost-effective musculoskeletal surgical procedures and with usage still growing worldwide [27], due to an increasing aging population and the expectation for higher functional capability in the later decades of life [28]. The most common reason for this type of procedure is arthritis, particularly when the knee does not respond to other treatments such as physical therapy or injections.

Osteoarthritis is a degenerative disease that affects the entire joint. It is characterized by an imbalance between the formation and breakdown of the cartilaginous matrix, leading to the loss of cartilage in localized areas. This increases focal stress on the joints, contributing to increasing cartilage loss. When degeneration reaches a certain extent, the joint may become tilted or malaligned[29–31]. Its progression varies from person to person, and the most common symptoms include pain, weakness, stiffness, and reduced mobility[29–31]. There is no cure, and treatment focuses on alleviating symptoms. Initial treatment includes physical therapy, pharmacotherapy,

and other methods to reduce knee load. When conservative treatment fails, joint-preserving surgical options may be considered, ranging from arthroscopic lavage, shaving, and debridement to autologous transplantation. If all of these treatments fail, then a joint replacement surgery is performed[29–31].

Rheumatoid arthritis is a chronic autoimmune condition in which the immune system attacks the body's own tissues, leading to progressive damage of synovial-lined joints such as the knee. This results in symptoms including inflammation, swelling, stiffness, and pain. The condition typically affects both knees simultaneously[32, 33]. Although significant advances have been made in the development of biological agents in the field of rheumatology, progressive joint destruction remains a problem in a subset of patients. For these individuals, TKA continues to be the most successful intervention for reducing symptoms[32, 33].

Post-traumatic arthritis can be seen as a type of osteoarthritis that develops after a joint injury, resulting from cartilage damage after trauma. It is more common to occur in more active individuals who participate more in activities that can cause knee injuries[34].

2.4 Total Knee Arthroplasty Surgery

The overarching objective of primary TKA is to alleviate persistent arthritic joint pain, correct pathological mechanical deformities, restore functional range of motion, and significantly improve patient-reported quality of life [35, 36]. To achieve optimal long-term prosthetic survivorship and high rates of postoperative satisfaction, the surgeon must balance two vital intraoperative objectives: executing precise lower limb alignment and establishing a symmetrical, well-balanced soft-tissue envelope throughout the entire arc of motion [35, 36].

2.4.1 Pre-operative Evaluation and Risk Stratification

Meticulous pre-operative planning and comprehensive clinical profiling are mandatory to mitigate perioperative complications and predict long-term functional recovery following elective TKA. Initial clinical investigations require a comprehensive physical examination to rule out coexisting hip or spine pathologies that might act as confounding pain generators, alongside an evaluation of prior surgical scars to prevent localized skin necrosis. Furthermore, vascular profiling via the palpation of distal pulses is necessary to screen for peripheral vascular disease, ensuring adequate baseline patency before scheduling surgery [35, 36].

Because patients scheduled for arthroplasty often present with complex systemic comorbidities, establishing target physiological optimization thresholds is critical for minimizing postoperative risk [36]. Clinical data demonstrate that stabilizing modifiable metabolic and physical variables prior to incision directly correlates with reduced lengths of hospital stay and a decreased incidence of superficial or deep wound complications [35, 36]. To achieve safe clinical clearance, the literature suggests optimizing high-risk patients to an acceptable body mass index and ensuring a robust baseline nutritional status prior to surgical intervention [35, 36].

Concurrently, systemic screening must target precise glycemic control and red blood cell parameters to avoid compromising the downstream tissue healing cascade. Poorly managed diabetes significantly escalates the risk of deep prosthetic joint infection, making strict pre-operative glycemic control the standard recommendation for surgical safety. Finally, a dedicated anemia workup should be initiated to ensure hemoglobin concentrations are sufficient to meet clinical safety parameters, effectively safeguarding the patient against independent drivers of perioperative morbidity and mortality [35, 36].

During pre-operative planning, radiographs are used to evaluate mechanical alignment, providing valuable information about existing deformities and bone loss. The femoral resection angle is calculated as the angle between the anatomical axis, defined by the intramedullary canal, and the mechanical axis of the femur, as selected by the surgeon, as shown in Figure 2.5. Alignment angles of approximately 6° have been associated with improved joint function and a lower perception of the artificial joint during daily activities. These angles are reproduced intraoperatively, with the distal femoral cut made perpendicular to the mechanical axis. Following appropriate soft-tissue balancing, the resected distal femoral and tibial surfaces should be parallel[5, 35, 36].

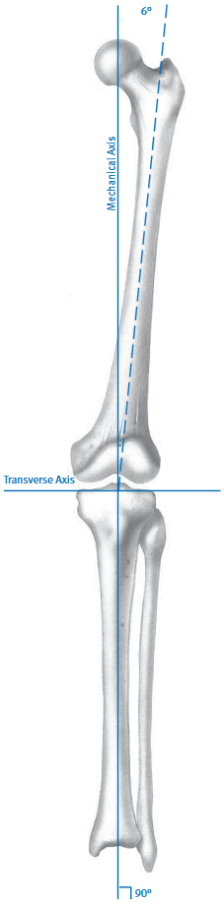
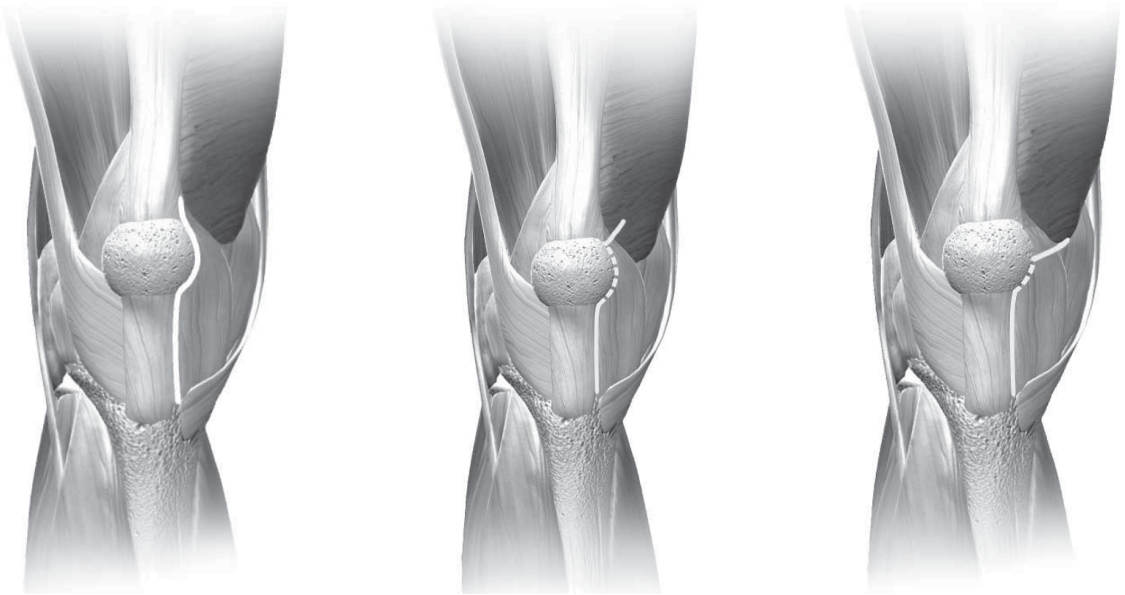


Figure 2.5: Femoral resection angle defined between the anatomical and mechanical axes of the femur [5]

2.4.2 Surgical Exposures and Approaches

Selecting an appropriate arthrotomy technique requires balancing the need for clear joint exposure against the desire to preserve the extensor mechanism and localized tissue vascularity [35, 36]. The primary approaches utilized in contemporary primary TKA are shown in Figure 2.6 and include:

- **Medial Parapatellar Approach:** This technique remains the historical benchmark and involves an incision through the medial aspect of the quadriceps tendon, continuing distally around the patellar cuff to elevate a continuous subperiosteal sleeve along the proximal medial tibia. It provides unparalleled surgical visualization for complex fixed deformities and is highly reliable when operating on obese populations, though violating the tendon can occasionally prompt early patellar tracking anomalies or transient quadriceps inhibition [35, 36].
- **Midvastus Approach:** Operating as a functional alternative, this muscle-splitting exposure completely spares the primary quadriceps tendon by separating the fibers of the VMO muscle belly obliquely toward the superomedial patellar border. This method strikes an effective clinical compromise, preserving early quadriceps function while offering adequate joint exposure for routine prosthetic implantation [36].
- **Subvastus Approach:** Characterized as a muscle-sparing technique, this approach avoids any direct incision into the quadriceps mechanism or the VMO muscle belly by cleanly lifting the lower boundary of the VMO off the intermuscular septum . While it preserves native patellar vascularity, its inherently restricted visibility increases technical difficulty in obesity patient profiles [36].

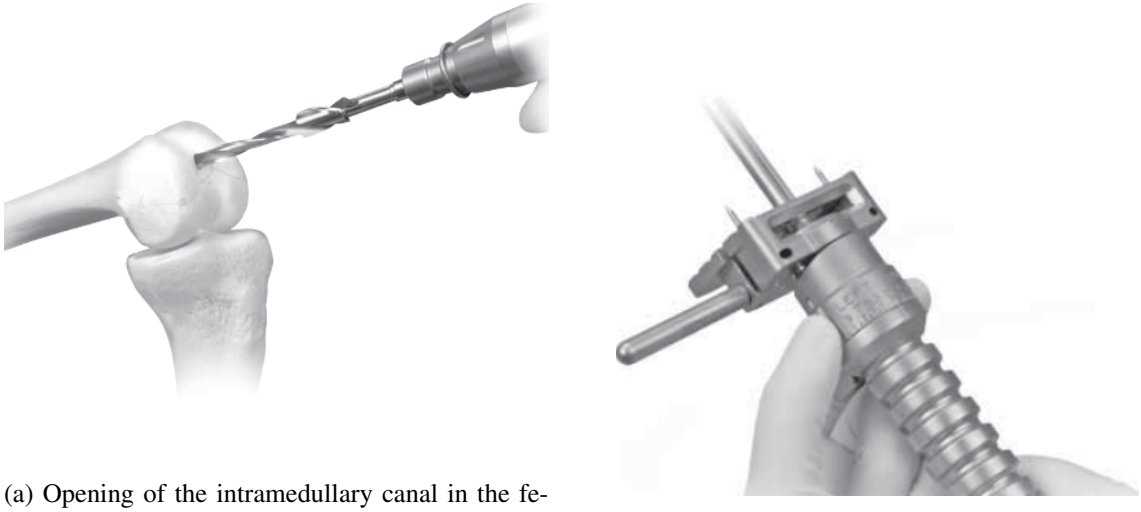


(a) Medial parapatellar approach (b) Midvastus approach (c) Subvastus approach

Figure 2.6: Alternative surgical approaches to the knee [5]

Primary TKA relies on a systematic sequence of structural bone resections combined with calculated soft-tissue balancing to achieve a stable, functional joint reconstruction, the order of the resections may vary depending on the surgeon preferences, here is presented an example approach [35, 36]:

– **Femoral Instrumentation and Resection:** Following patellar eversion and joint displacement, the femoral canal is opened with an intramedullary drill to accommodate an alignment guide, as shown in Figure 2.7. The guide is adjusted to the patient's native anatomy, typically set to 5° or 7° of valgus as determined by the pre-operative radiographs, to direct a distal femoral bone cut of approximately 9 to 10 mm, which matches the thickness of the prospective femoral implant [35, 36].

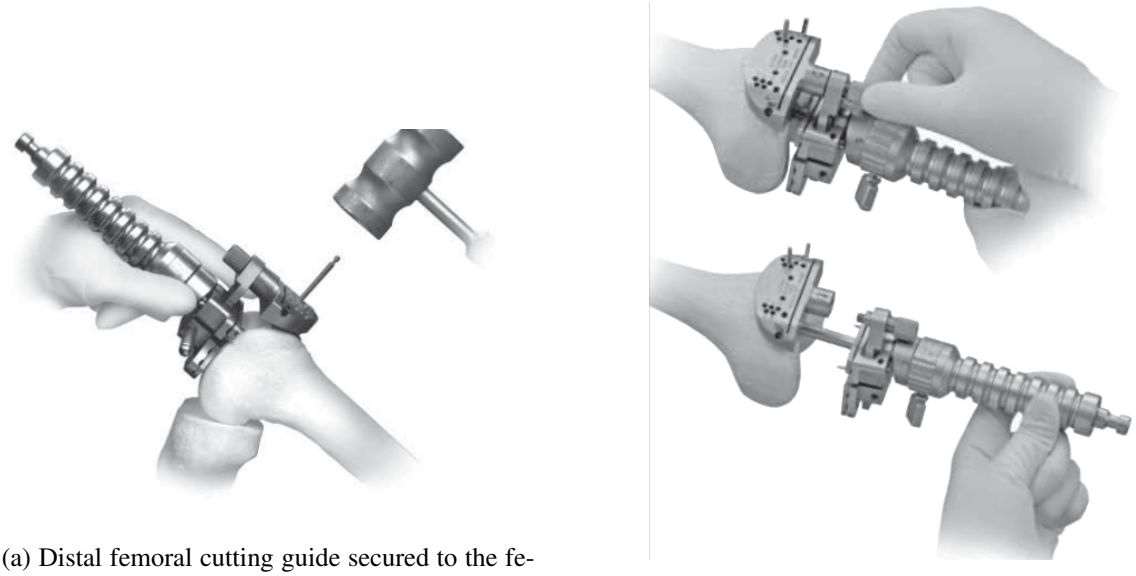


(a) Opening of the intramedullary canal in the femur using a drill

(b) Alignment guide

Figure 2.7: Initial steps of the femoral preparation [5]

The alignment guide is inserted into the intramedullary canal of the distal femur, and the distal femoral cutting guide is attached to it. After this is still possible to do minor adjustments made as necessary to ensure proper positioning. Once the desired position is achieved, the cutting guide is securely fixed to the femur, and the alignment guide is removed, as shown in Figure 2.8. Subsequently, a rotating saw is inserted through the cutting slot of the guide to perform the distal femoral resection[5], as shown in Figure 2.9.



(a) Distal femoral cutting guide secured to the femur

(b) Removal of the alignment guide

Figure 2.8: Positioning and securing of the distal femoral cutting guide [5]



Figure 2.9: Distal femoral resection using a rotating saw [5]

After the distal femoral cut, the knee is flexed to 90°, and a sizing guide is attached to the resulting flat surface, as shown in Figure 2.10. The sizing guide is used to measure the femur and determine the appropriate size of the 4-in-1 cutting guide. The selected 4-in-1 cutting guide is then positioned on the distal femur and securely fixed in place. A saw is subsequently used to perform the four required femoral cuts[5], as shown in Figure 2.11b.

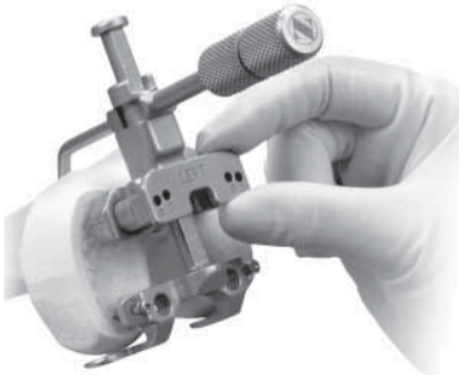
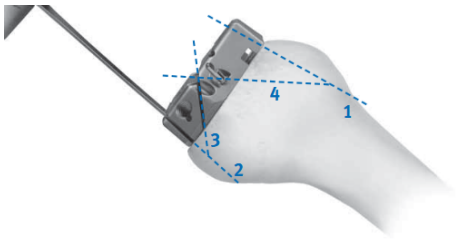
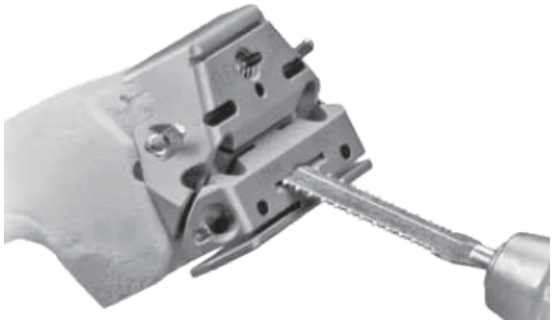


Figure 2.10: Sizing guide attached to the distal femur [5]



(a) Reference cuts of the 4-in-1 cutting guide



(b) A femoral cut using the 4-in-1 cutting guide

Figure 2.11: Positioning of the 4-in-1 cutting guide and execution of the femoral resections [5]

– **Tibial Alignment and Resection:** The proximal tibia is prepared using either an extramedullary or intramedullary surgical guide, as shown in Figure 2.12. The cut is typically oriented perpendicular to the mechanical axis of the tibia, though surgeons pursuing specialized anatomic alignment protocols may introduce a conservative 2° to 3° varus slope. Accurate rotational alignment of the tibial baseplate is secured by referencing the medial one-third of the tibial tubercle, cross-checking with the second ray of the foot, and tracking along the anatomical tibial crest [35, 36].



Figure 2.12: Tibia cutting guide [5]

– **Gap Evaluation and Dynamic Balancing:** Once the distal femoral and proximal tibial bone cuts have been completed, the knee is assessed in both extension and flexion using a structural spacer block to evaluate the resulting gaps under appropriate tension. With the knee in full extension, the spacer block is inserted into the joint space and an external alignment rod is used to verify that the extension gap is balanced and rectangular, and that the coronal mechanical axis has been appropriately restored [35, 36]. The knee is then flexed to approximately 90° to assess the flexion gap and the rotational alignment of the femoral component. The relative tension between the extension and flexion gaps is then evaluated to determine whether further bone resection or soft-tissue adjustment is required[35, 36].

– **Trial Reduction and Patellofemoral Preparation:** Trial components are positioned to evaluate joint stability throughout a full range of motion, confirming that the knee remains balanced from 0° of extension into deep flexion . If patellar resurfacing is indicated, the bone must be resected symmetrically to prevent patellofemoral overstuffing, which causes chronic anterior knee pain, or excessive bone loss, which risks post-operative fracture. Once patellar tracking is confirmed, the definitive components are securely impacted [35, 36].

Chapter 3

State of Art

3.1 Point Cloud Registration

To register a patient's anatomy in real time, the system must first capture the geometry of the exposed bone surfaces intraoperatively. Traditional computer-assisted surgery systems rely on mechanical point tracing using optical styluses or physical probes to manually sample discrete anatomical landmarks, which creates a sparse point cloud and introduces operational latency [37]. Conversely, modern approaches leverage non-contact optical sensors to facilitate continuous, markerless surface registration. By combining real-time depth cameras with automated segmentation, these systems directly align the exposed surgical field with pre-operative anatomical models using iterative point cloud registration methods [38].

Table 3.1 provides a concise summary of the local refinement, global coarse alignment, and deep learning-based registration frameworks discussed in this section.

3.1.1 Rigid Point Cloud Registration Algorithms

The fundamental objective of rigid 3D point cloud registration in computer-assisted orthopaedic surgery is to establish an accurate spatial alignment between intraoperatively acquired patient anatomy and a high-resolution pre-operative surgical plan. Mathematically, this task consists of estimating the optimal rigid transformation parameters, a 3D rotation matrix and translation vector, that map the moving source point cloud onto the fixed target model [39]. This is typically solved by minimizing an objective function that quantifies the residual error between corresponding geometric features.

3.1.1.1 Local Refinement Frameworks

The Iterative Closest Point (ICP) algorithm stands as the foundational standard for local alignment in medical point cloud registration, iteratively establishing explicit nearest neighbour point correspondences between source and target geometries to minimize the spatial distance between them [39, 40]. Despite its widespread use, classic point-to-point ICP [41] exhibits major limitations in surgical environments, because it forces points to align precisely with immediate geometric neighbors without accounting for underlying surface contours, it is highly sensitive to the initial pose, sensor noise, and outliers [40]. Conversely, point-to-plane ICP minimizes distances along target surface normals, allowing points to slide along smooth anatomical surfaces and improving convergence stability [42, 43]. Nonetheless, narrow convergence profiles can still trap local algorithms in incorrect minima if the initial setup is sub-optimal [39, 40, 44].

To handle the incomplete spatial overlap typical of intraoperative surface scanning, variations like Trimmed ICP are deployed to discard outlier regions [40, 45]. Rather than optimizing over the entire point set, these frameworks utilize a robust estimation approach that computes objective metrics only on a designated subset of high-confidence matches, allowing the system to disregard non-overlapping regions [40, 45]. Furthermore, real-world orthopaedic settings introduce severe structural noise, including soft tissue displacement, blood, surgical fluids, and non-uniform sensor sampling. To manage these artifacts, advanced local refinement pipelines transition from strict deterministic matching to robust probabilistic formulations, density-adaptive methods like DARE jointly infer the latent scene distribution and transformation parameters by minimizing their Kullback-Leibler divergence, maintaining robustness against severe sampling and viewing-angle variations common to depth sensors[46].

Additionally, batch frameworks like joint registration of multiple point clouds (JRMPC) formulate multi-view registration as a generative Gaussian Mixture Model, jointly estimating transformations against a shared latent shape to eliminate pairwise error accumulation and reference-frame bias [47]. Finally, filter-based frameworks like FilterReg reformulate a probabilistic expectation step as a fast Gaussian filtering task paired with twist parameterization, surpassing standard ICP in computational speed while maintaining high resilience to sensor noise and partial occlusions[48]. To overcome the sensitivity of discrete point-to-point correspondences, continuous probabilistic and cell-based local formulations are also deployed. Coherent point drift treats registration as a probability density estimation problem, forcing moving source points to deform or translate coherently as a Gaussian Mixture Model [49]. Similarly, the 3D Normal Distributions Transform (NDT) partitions the target space into a grid of 3D cells, modeling the local geometry of each cell as a continuous probability density function defined by a local mean and covariance matrix [50]. By optimizing alignment directly against this continuous distribution field via numerical optimization,

3D-NDT eliminates explicit nearest-neighbor lookups while remaining resilient to sensor density variations

3.1.1.2 Global and Coarse Alignment Frameworks

To operate reliably in a clinical environment without demanding precise manual positioning from the surgeon, modern systems implement automated global alignment frameworks. Global approaches overcome the initialization dependency inherent to local algorithms but often at the expense of computational cost [51], their primary purpose is to achieve an acceptable baseline alignment without requiring precise manual initialization or when the initial poses are completely unaligned, scanning the global spatial properties of the anatomy to prevent subsequent local refinement loops from falling into local extrema [39].

Traditional automated pipelines implement a two-stage, feature-based matching strategy. First, hand-crafted geometric descriptors, most notably Fast Point Feature Histograms (FPFH), are extracted to capture multidimensional angular and distance relationships between surface normals within local neighborhoods [52]. Second, a Random Sample Consensus (RANSAC) framework pairs these descriptors, randomly sampling minimal point subsets to generate rigid transformation hypotheses [53]. The transformation that maximizes the global consensus (inlier count) is selected, providing a robust baseline alignment even under extreme initial displacements. Alternatively, Fast global registration pairs these descriptors without iterative RANSAC sampling by directly optimizing an objective function across candidate matches via graduated non-convexity, achieving faster execution while maintaining robustness to false correspondences [54].

For guaranteed global optimality, Branch-and-Bound frameworks recursively partition the $SE(3)$ transformation space, evaluating upper and lower objective bounds to discard suboptimal regions. In practice, Go-ICP computes these bounds directly on Euclidean point distances [40], whereas GOGMA applies them to continuous gaussian mixture models to bypass discrete correspondence searches entirely [55].

3.1.2 Deep Learning-Based Registration Paradigms

To overcome the inherent limits of hand-crafted local descriptors, recent computer-assisted orthopaedic pipelines transition toward deep data-driven feature extraction directly from unstructured, raw 3D point cloud distributions. PointNet pioneered this paradigm by establishing a permutation-invariant framework that passes spatial coordinates through shared multi-layer perceptrons and aggregates them via a symmetric max-pooling function [56]. However, because PointNet computes a purely global feature representation, it fails to capture local geometric interactions. To overcome this limitation and resolve subtle anatomical

variations across bone surfaces, PointNet++ introduces a hierarchical structure that recursively partitions point sets into localized neighborhoods to extract multi-scale geometric features[39].

Building on these point-based approaches, subsequent feature extraction strategies organize unstructured point clouds using spatial grids, voxels, or local neighbourhoods to capture richer geometric relationships. Convolutional operations and learnable kernels can then extract local and multi-scale features from these structured representations, while local reference frames can further improve rotation invariance. Such strategies enable more distinctive and compact descriptors for efficient correspondence matching, although they may introduce additional computational costs and remain sensitive to noise, normal estimation, initial pose, and variations in point cloud density [39].

Beyond feature extraction, deep learning-based registration methods also address keypoint detection, outlier rejection, and motion estimation. Keypoint detection aims to identify informative points that are less affected by noise and variations in point-cloud density, while outlier rejection removes unreliable correspondences caused by noise, partial overlap, and insufficient local features. Once reliable correspondences are established, the rigid transformation can be estimated using approaches such as single value decomposition based optimization or learned regression of the motion parameters [39].

3.1.3 Point Cloud Registration in Total Knee Arthroplasty

TKA demands highly specialized point cloud registration pipelines to achieve sub-millimeter and sub-degree alignment of two distinct, articulating structures: the distal femur and the proximal tibia. Because registration errors directly propagate to misaligned component positioning, altered joint kinematics, and premature implant wear, maintaining high accuracy is critical[57]. However, the limited intraoperative surgical window provides only a partial, unilateral view of the femoral condyles and tibial plateau, resulting in an incomplete source point cloud when compared to full-volume pre-operative computed tomography(CT) or magnetic resonance imaging (MRI) target models[58, 59].

Furthermore, the physical environment of an arthritic knee joint introduces structural noise that complicates standard geometric alignment:

- **Osteophytic Formations:** Pathological bone overgrowths alter the surface topology of the peripheral femur and tibia. If captured intraoperatively but poorly segmented in pre-operative models, these structures act as geometric outliers that severely bias local refinement loops like ICP [60].
- **Articular Cartilage Degradation:** Non-uniform cartilage wear exposes subchondral bone, causing a structural mismatch because optical surface scanners capture the superficial boundary (cartilage/exposed bone) while pre-operative CT models delineate purely cortical bone [61].

– **Soft Tissue Interferences:** Remnants of the meniscus, cruciate ligaments, and surrounding joint capsule inject non-rigid structural artifacts into the rigid point cloud distribution [62].

To mitigate these challenges, state-of-the-art TKA registration systems employ a multi-stage, region-of-interest weighted alignment strategy[63]. During the intraoperative surface acquisition phase, advanced software pipelines dynamically segment the point cloud, assigning higher statistical weights to stable cortical regions, such as the anterior femoral cortex and the non-eroded areas of the condyles while down-weighting or masking out osteophyte-prone boundaries[63]. Alternatively, recent algorithmic developments have sought to minimize sensitivity to poor sampling and segmentation noise at the distal femur by reversing the registration direction within a Bounded-ICP framework [62]. To address partial anatomical exposure, point sparsity, and intraoperative noise, partial-to-full registration can be formulated using a continuous Gradient Signed Distance Field [59]. By optimizing intraoperative points directly against the pre-operative distance field via a robust Cauchy M-estimator, this approach eliminates the need for explicit correspondences or manual landmark selection, enabling stable convergence from sparse, randomly probed bone surfaces [59].

In an image free platform like the CT-free VectorVision navigation system, two passive reference arrays are rigidly attached to the femur and tibia, while an optical tracking system is used during registration. The registration workflow defines the hip centre of rotation using a pivoting procedure, digitizes the medial and lateral malleoli and specific knee landmarks, and records bone-surface information by sliding a navigation pointer over the tibial plateau and femoral condyles. Based on these data, the system creates an adapted bone model of the knee[64].

Recent deep learning frameworks for TKA leverage point cloud completion architectures to recover unexposed anatomical geometry before surface registration. For example, methods utilizing Point Completion Networks predict the full femoral shape from minimal intraoperative data, including digitized distal points and the femoral head center [65]. Registering this reconstructed full bone via iterative closest point avoids local minima and resolves rotational ambiguities, ensuring the global femoral pose accuracy required across diverse surgical alignment strategies[65].

3.2 Projection Mapping

Table 3.2 synthesizes the static calibration frameworks and dynamic projection mapping approaches discussed in this section, summarizing their operating principles and specific characteristics.

3.2.1 Static Projector-Camera Calibration Fundamentals

In SAR and structured-light pipelines, geometric calibration establishes the mathematical mapping between two-dimensional (2D) image coordinates and the 3D physical workspace. A central principle in these systems is the projector-camera duality where the projector is mathematically modeled as an inverse pinhole camera, projecting light rays outward from its spatial light modulator rather than capturing incoming rays on a sensor.

Consequently, both devices share the same projective geometry. The mapping of a homogeneous 3D world point $P = [X, Y, Z, 1]^T$ to a homogeneous 2D pixel coordinate $q = [u, v, 1]^T$ is expressed up to a scale factor s by:

$$sq = K[R \mid t]P \tag{3.1}$$

where K is the 3×3 intrinsic matrix encoding the internal geometry of the device (camera K_c or projector K_p), and $[R \mid t]$ is the 3×4 extrinsic matrix defining the device's spatial pose relative to the world coordinate frame.

In real-world optical systems, physical lenses introduce systematic geometric deviations from this idealized model. These are categorized into radial distortions, caused by lens curvature variations, and tangential distortions, resulting from misalignment between the lens elements and the image sensor plane [66]. To recover true spatial coordinates, literature relies on estimating distortion coefficients to mathematically rectify the captured imagery.

For system calibration, Zhang's foundational framework [66] revolutionized the field by replacing complex 3D rigs with a easily manufactured 2D planar target (e.g., a checkerboard). By observing the planar target across multiple orientations, the 3D-to-2D projection collapses into a series of 2D-to-2D homographies, allowing the explicit decoupling and calculation of both intrinsic and extrinsic parameters.

To integrate projection capabilities into this framework, modern state-of-the-art architectures model the projector as an inverse pinhole camera that projectively emits light rays rather than capturing them . By adapting Zhang's planar homography approach, structured-light patterns are mapped onto the calibration target, enabling the joint estimation of the projector's intrinsic parameters and its relative extrinsic transformation matrix ($[R_{cam}^{proj} \mid t_{cam}^{proj}]$) to establish a unified triangulation framework [67–70].

Alternative static calibration frameworks vary based on the required level of automation and the sensor configuration. In some interactive alignment systems, predetermined 3D computer-aided design (CAD) models are aligned with physical objects through manual calibration, utilizing software to adjust the virtual model to its physical counterpart[71]. To eliminate manual intervention, automated self-calibration frameworks for multi-projector-camera systems establish correspondences without prior knowledge of device parameters.

This is achieved either by projecting structured light patterns to initialize a geometric reconstruction from a baseline camera pair incrementally adding devices via non-linear optimization [72] or by projecting patterns onto a translucent sheet [73], because the features are visible from both sides of the sheet, multiple cameras capture the same patterns to resolve joint projector-camera and camera-camera intrinsic and extrinsic parameters [73]. Additionally, to account for non-planar surfaces and long projection distances, setups incorporate projector and light detection and ranging (LiDAR) calibration. This approach estimates the system parameters to map 3D LiDAR points directly to 2D projector pixels, ensuring the projected image warps to match the real-world geometry [74].

3.2.2 Dynamic Projection Mapping

The progression from static to dynamic projection mapping (DPM) is driven by advancements in processing architectures and the availability of high-frame-rate cameras and projectors. DPM frameworks are structurally categorized based on the underlying surface kinematics. Rigid transformations are characterized by a global geometric relationship where the target object moves relative to the projection system without surface deformation. Conversely, non-rigid transformations involve deformable surfaces, requiring the projected imagery to adapt dynamically to localized topological variations [75, 76].

A primary operational challenge in DPM is the management of occlusions, which occur when the optical path between the hardware components and the target surface is obstructed. Occlusions are classified as external when an independent object blocks the line of sight, or as self-occlusions when a local deformation or geometric orientation of the target surface shields a region from the sensor or projector field of view [75, 76].

3.2.3 Rigid Projections

DPM for rigid transformations typically demands less computational overhead than non-rigid frameworks, as the pipeline prioritizes rigid pose estimation over dense, time-varying surface reconstructions. Implementations are broadly divided into marker-based and markerless tracking architectures.

For marker-based approaches, embedded tracking configurations can minimize visual interference. One methodology integrates physical markers within the internal structure of the object, optimizing their spatial distribution to maximize the number of valid viewpoints [77]. This architecture utilizes a stereo infrared (IR) camera setup to track the markers independently of the visible-light projection, employing radiometric compensation algorithms to reduce the surface visibility of the tracking elements. While effective across varying surface textures, the approach is constrained to single-object scenarios and exhibits processing latencies that limit real-time interactivity.

To eliminate object modification, markerless frameworks leverage geometric features. A dual-camera IR tracking pipeline utilizes Canny edge detection to extract the external contours of the target object [78]. Operating in the IR spectrum isolates the tracking sequence from the visible projection, while the integration of multi-view contour data resolves spatial ambiguities and filters out external occlusions by discarding correspondence points outside the predicted object boundary.

Alternatively, low-latency rigid tracking can be achieved by combining multi-spectral projection and temporal filtering. A representative architecture pairs an IR projector, an IR camera, and a visible-light RGB projector [79]. The system executes high-speed rigid-body tracking by projecting structured-light profiles, consisting of Gray code and phase-shifted sinusoidal patterns, within the IR spectrum. The synchronized IR camera captures these patterns to continuously update surface depth maps through the temporal integration of successive frames, estimating object motion and geometry under low-latency constraints.

3.2.4 Non-Rigid Projections

Non-rigid projection mapping requires tracking continuous surface deformations over time. Standard tracking strategies often rely on surface-mounted marker arrays to capture localized displacements.

Low-latency tracking of deformable surfaces can be achieved using high-speed cameras synchronized with dot marker arrays [80]. This framework mitigates temporary marker occlusions by employing fast temporal tracking algorithms, maintaining stable estimations of the underlying surface deformation during intermittent data loss. To precisely define the projection boundaries on non-rigid targets, retro-reflective dot marker arrays can be paired with an auxiliary depth camera [81]. The depth data delineates the peripheral boundaries of the deformable object, aligning the projected content with the moving edges of the target while preventing projection spillover onto surrounding surfaces.

To prevent the visual clutter associated with standard physical markers, frameworks utilize invisible tracking elements. Applying IR-absorbing ink to create surface textures allows an IR camera to monitor geometric changes without compromising the visual appearance of the projection [82]. The localized tangential deformation is estimated by calculating the displacement vectors of the IR markers between consecutive frames. Furthermore, machine learning architectures can enhance prediction accuracy, integrating retro-reflective markers with multi-camera motion capture systems and deep neural networks enables the real-time prediction of the 3D position and pose of the deforming surface, allowing the rendering pipeline to pre-warp imagery based on anticipated geometric states [83].

When physical modifications to the target surface are impractical, such as in applications involving biological tissues or fluids, markerless non-rigid tracking is required. Texture-based markerless methods estimate the non-rigid warp between successive frames by analyzing variations in pixel intensity across the surface topology [84]. These algorithms incorporate self-occlusion reasoning to automatically detect obscured regions, systematically down-weighting these zones within the optimization cost function to prevent incorrect geometric correspondences[84].

Table 3.1: Summary of local refinement, global alignment, and deep learning-based point cloud registration methods

Category	Method	Principle	Key Characteristics
Local Refinement	Point-to-Point ICP [41]	Nearest-neighbor distance minimization	Simple baseline; sensitive to initialization, noise, and outliers
	Point-to-Plane ICP [42]	Normal-projected distance minimization	Better stability on smooth surfaces; still prone to local minima
	Trimmed ICP [45]	Optimization over selected matches	Handles incomplete overlap by rejecting outlier regions
	DARE [46]	KL-divergence between scene distributions	Density-adaptive; robust to sampling and viewing-angle variations
	JRMPC [47]	GMM-based joint estimation	Multi-view registration using a shared latent shape.
	FilterReg [48]	Gaussian filtering with twist parameterization	Faster than ICP; robust to noise and partial occlusions
	Coherent point drift [49]	Probabilistic GMM formulation	Continuous formulation; coherent source-point motion
	3D-NDT [50]	Grid-based continuous probability model	Avoids nearest-neighbor searches; handles density variations
Global & Coarse	FPFH + RANSAC [52, 53]	Feature matching with random consensus sampling	Robust to large initial displacement
	Fast Global Registration [54]	Graduated non-convexity optimization	Faster alternative to RANSAC-based matching
	Go-ICP [40]	Branch-and-Bound over $SE(3)$	Globally optimal; avoids local-minimum entrapment
	GOGMA [55]	Branch-and-Bound with GMM bounds	Globally optimal; avoids discrete correspondences
Deep Learning	PointNet [56]	Shared MLPs and max pooling	Permutation-invariant; lacks local geometric interactions
	PointNet++ [39]	Hierarchical local neighborhoods	Captures multi-scale local features
	Structured Representations [39]	Grid, voxel, and neighborhood features	Richer multi-scale descriptors; can improve rotation invariance
	Deep Pipeline Modules [39]	Keypoint, outlier, and motion networks	Supports correspondence filtering and rigid pose estimation

Table 3.2: Summary of static calibration and dynamic projection mapping methods discussed in the literature

Category	Approach & References	Working Principle	Key Characteristics
Static Calibration	Manual CAD Alignment [71]	Software-based adjustment to align pre-existing 3D CAD models with physical objects	Requires manual intervention
	Planar Homography & SL [67–70]	Projects structured-light patterns onto a 2D planar target to estimate K and $[R_{cam}^{proj} \mid t_{cam}^{proj}]$	Decouples intrinsic and extrinsic parameters into a unified triangulation model
	Automated Multi-Device [72, 73]	Incremental optimization from baseline pairs or dual-sided capture on a translucent sheet	Eliminates manual intervention; resolves joint parameters across multiple devices
	Projector–LiDAR [74]	Direct mathematical mapping from 3D LiDAR points to 2D projector pixels	Adapts projection warping to non-planar surfaces and long projection distances
Rigid Dynamic Projection	Embedded IR Markers [77]	Internal markers tracked via stereo IR cameras with radiometric compensation	Reduces visual interference; constrained to single objects and subject to latency
	Markerless Contour [78]	Dual-camera IR Canny edge extraction with multi-view boundary validation	No object modification; discards external points to filter occlusions
	Multispectral SL [79]	IR projector (Gray code + phase shift), IR camera, and visible-light RGB projector	Temporal integration of depth maps for low-latency rigid tracking
Non-Rigid Dynamic Projection	High-Speed Dot Arrays [80, 81]	Synchronized high-speed cameras or retro-reflective markers with auxiliary depth sensing	Fast temporal tracking handles intermittent loss; depth prevents boundary spillover
	IR-Absorbing Ink [82]	Invisible surface ink patterns tracked via IR camera displacement vectors	Eliminates visual marker clutter; estimates tangential deformations
	MoCap with DNN [83]	Retro-reflective markers with multi-camera motion capture and neural networks	Predicts real-time 3D pose to pre-warp imagery for anticipated geometric states.
	Markerless Texture [84]	Surface pixel intensity analysis with self-occlusion cost down-weighting	Non-invasive for tissues and fluids; mitigates errors from obscured zones

Chapter 4

Methodology

This chapter describes the methodology developed to implement a SAR system utilizing a setup composed by a projector and a depth sensor. The primary objective of the system is to assist distal femoral resection by determining and projecting the target anatomical position for a distal femoral cutting guide. This resection was selected as the primary focus because it offers the direct clinical benefit of avoiding intramedullary drilling and bypassing conventional mechanical alignment guides, thereby reducing the need for additional surgical instrumentation.

4.1 Development Workflow

The development of the solution is structured into four main components:

1. **Surgical Planning Application:** A software module was developed to enable the definition of the surgical resection plane on a 3D anatomical model of the patient's femur. This application computes the rigid body transformation between a specific virtual surgical guide and the patient's femur model for the planned cut configuration.
2. **3D Point Cloud Registration:** A spatial registration pipeline designed to align the coordinate system of the pre-operative 3D femur point cloud with the real-time point cloud acquired by the depth sensor, thereby localizing the physical femur within the system's workspace.
3. **Projector-Camera Calibration:** A calibration procedure executed to estimate both the intrinsic and extrinsic parameters of the projector and depth camera, establishing a common coordinate framework.
4. **Spatial Augmented Reality Projection:** A visualization pipeline that employs the computed transformation matrices to project the outer border of the surgical guide's

upper surface directly onto the physical environment, enabling the accurate positioning of the physical instrument according to the pre-operative surgical plan.

4.2 System Workflow

The overall operational workflow of the proposed system is outlined in Figure 4.1, divided into two primary stages: the Pre-Operative Planning phase and the Intraoperative phase. One-time procedures, such as projector-camera calibration, are performed beforehand and are therefore excluded from the routine surgical pipeline.

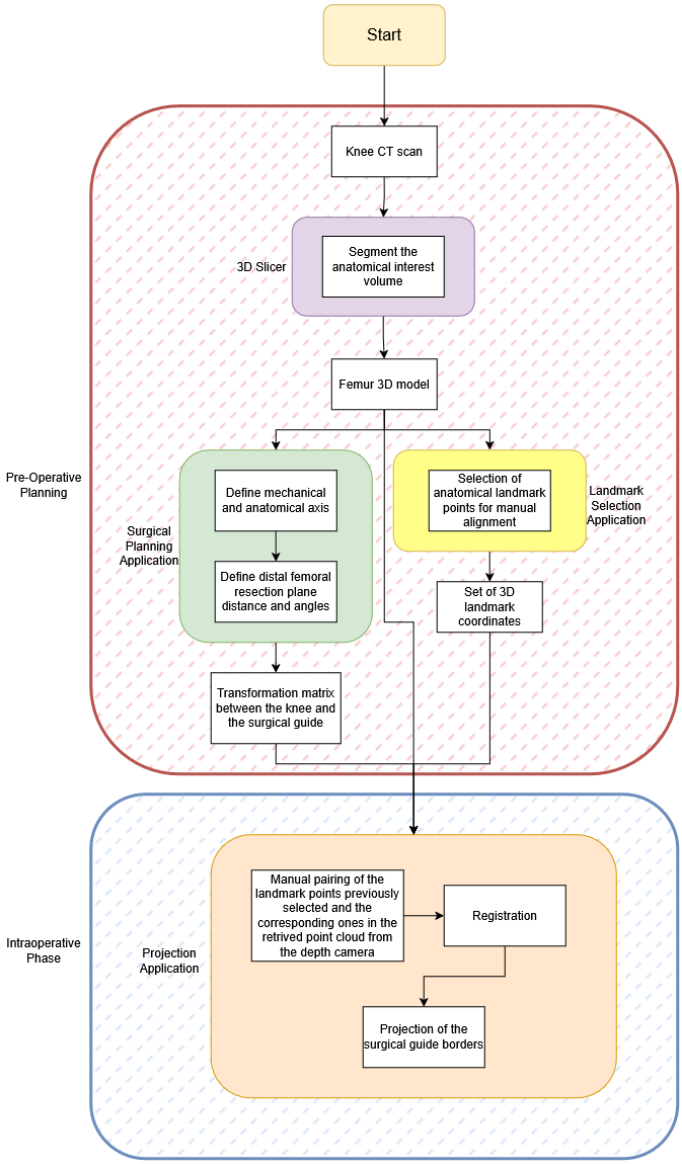


Figure 4.1: End-to-end operational workflow of the navigation system, encompassing pre-operative planning and intraoperative registration and projection

4.2.1 Equipment, Software and Tools

The physical experimental setup comprised an Optoma EW330 projector, an Ensenso N36 3D stereo camera, and a custom-configured computer system equiped with an AMD Ryzen 3 3100 CPU, 16 GB of DDR4 RAM, an NVIDIA GeForce RTX 3060 GPU with all components being installed on an MSI A320M-A Pro motherboard, providing a reliable platform for point cloud processing.

The pipeline was implemented using a combination of dedicated software programs and Python libraries. Within the Python environment, Qt6 [85] was used to build the graphical user interface, VTK [86] handled real-time 3D rendering and mesh manipulation, Open3D [87] performed point cloud related processing and surface registration, and OpenCV [88] executed computer vision tasks, image processing, and calibration routines. For external modeling and hardware setup, 3D Slicer [89] was employed to segment CT scans into 3D bone models, SolidWorks [90] was used for CAD of the cutting guide and physical fixtures, and Calib.io [91] was used to generate calibration targets for the projector-camera setup.

4.3 Surgical Planning Application

A pre-operative surgical planning application was developed to determine the optimal spatial positioning of the surgical cutting guide and to define the trajectory of the distal femoral resection. The application's graphical user interface and 3D visualization engine were implemented in Python, using the Qt6 framework for interface management and VTK for 3D mesh visualization. The layout architecture was designed and structured using Qt Designer.

The interface configuration is shown in Figure 4.2. The workspace is divided into distinct functional areas. The primary area consists of a central 3D visualization window for rendering and manipulating the 3D anatomical femur model.

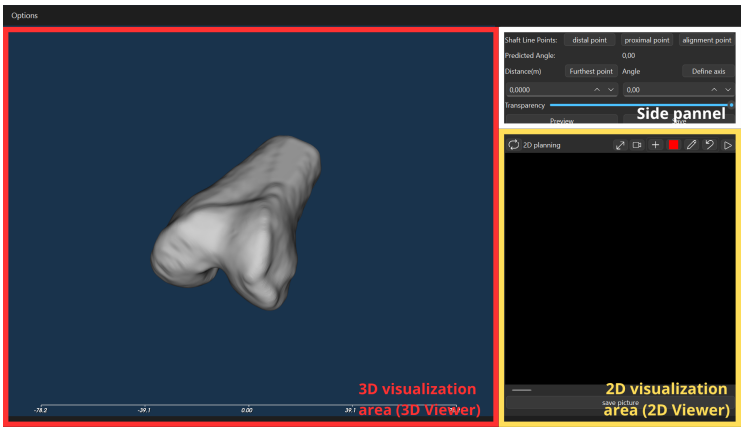


Figure 4.2: Surgical planning application user interface layout showing the 3D visualization area (red), 2D visualization area (yellow), and side panel (white)

Adjacent to this main window, a side panel consolidates the planning tools and parameters. This panel contains controls for defining anatomical reference points, adjusting planning parameters such as resection depth and alignment angles, and modifying visualization settings like model opacity.

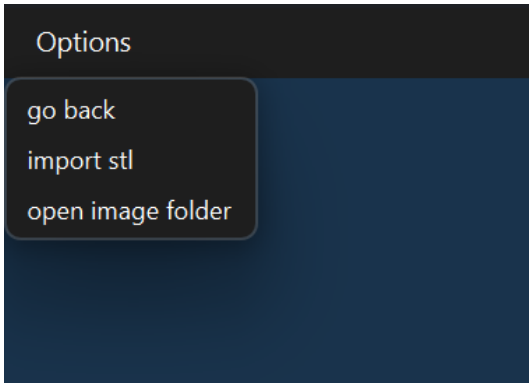
The interface also includes a secondary 2D viewer module that can display either medical images imported from local directories or dynamically generated upper and lower views of the bone slice at the user-defined resection plane. Additionally, a toolbar provides access to utility functions, including switching visualization modes, adding spatial annotations, taking measurements, and controlling the camera position.

4.4 Surgical Planning Pipeline

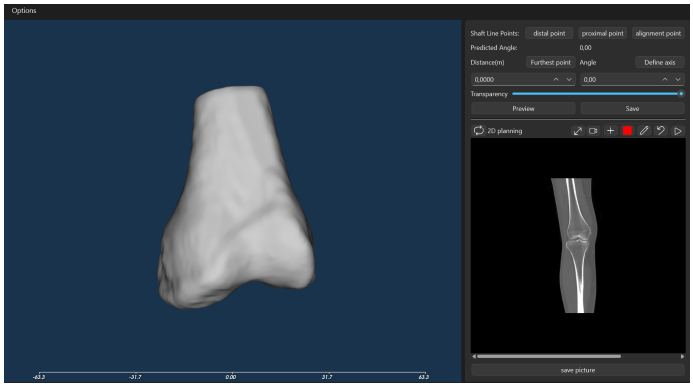
4.4.1 Femur 3D model

To initiate the surgical planning process, it is first necessary to obtain an stereolithography (STL) model of the femur bone. For this purpose, the open-source third-party software 3D Slicer was used together with the TotalSegmentator module [92], which enables the automatic segmentation of the femur from a knee CT scan. This approach significantly simplifies the generation of patient-specific anatomical models while reducing the manual effort typically associated with medical image segmentation.

After the segmentation process, the generated STL model is imported into the surgical planning application through the toolbar options shown in Figure 4.3a. Once imported, the 3D femur model can be visualized and manipulated within the main visualization area of the interface, allowing the healthcare professional to analyze the anatomical structure from different perspectives during the planning process.



(a) Toolbar options available in the surgical planning interface



(b) Surgical planning interface after importing the femur model and the image folder

Figure 4.3: Surgical planning interface and available toolbar options

To support the planning procedure, a directory import function was implemented to load medical images from a selected folder into the 2D viewer. These images provide complementary clinical data to assist in the planning and decision making process. Navigating through the imported dataset is achieved via an horizontal scroll bar located below the 2D viewport.

4.4.2 Anatomical Landmark Definition and Axis Alignment

To establish the anatomical axis within the planning application, the user manually defines two reference points based on previous imagiology studies: a distal point and a proximal point. The line intersecting these two points corresponds the intermedullary shaft of the femur, as shown in Figure 4.4.

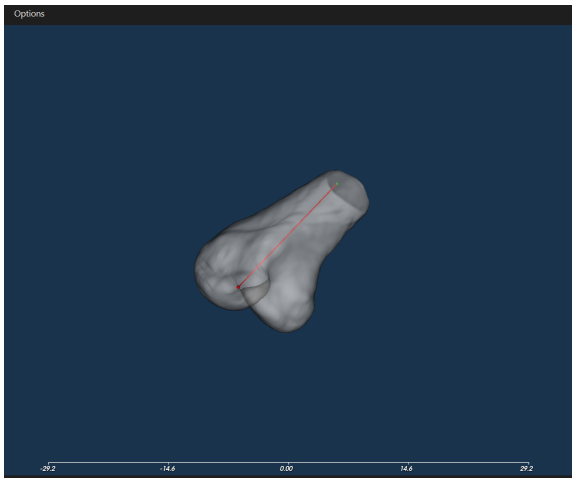


Figure 4.4: Definition of the femoral anatomical axis (red) via proximal and distal landmark selection

In conventional TKA, mechanical instrumentation relies on the angle between the anatomical and mechanical axis to configure in the alignment guide the position for the femoral cutting guide. To replicate this constraint virtually, the application utilizes the defined distal point alongside an additional entry point, represented as alignment point in the application, selected along the femoral shaft to establish the mechanical axis, as shown in Figure 4.5.

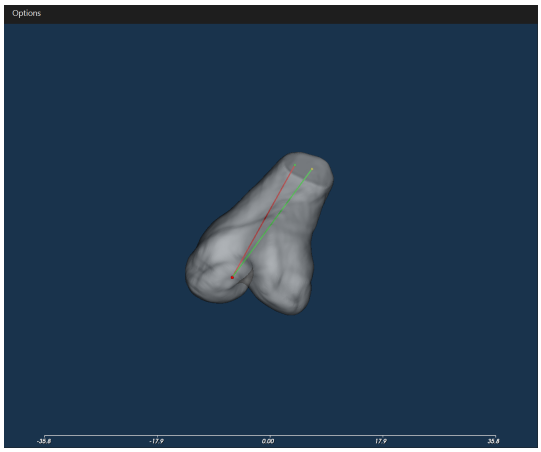


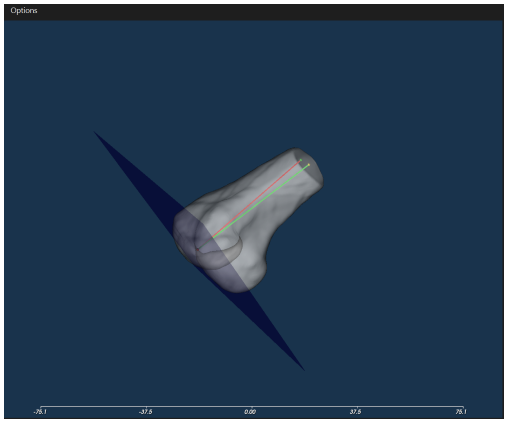
Figure 4.5: Simultaneous visualization of the virtual mechanical axis (green) and the anatomical axis (red)

By calculating the angular deviation between these two vectors, the software determines the femoral mechanical-anatomical angle, as shown in Figure 4.6.

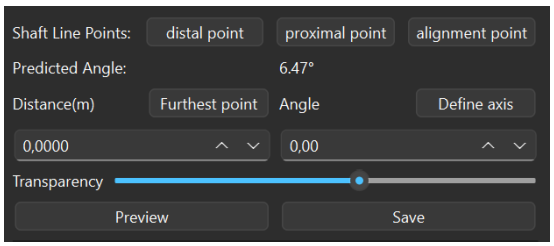


Figure 4.6: Calculated angular deviation between the mechanical and intramedullary axes, representing the projected distal cut angle

Upon establishing the anatomical axis, an initial distal cutting plane is automatically initialized. By default, this plane is constructed to be strictly perpendicular to the computed axis, with its coordinate origin positioned at the previously defined distal point, as shown in Figures 4.7.



(a) Cutting plane origin definition in the 3D planning environment (dark blue) with simultaneous visualization of the virtual mechanical axis (green) and the anatomical axis (red)

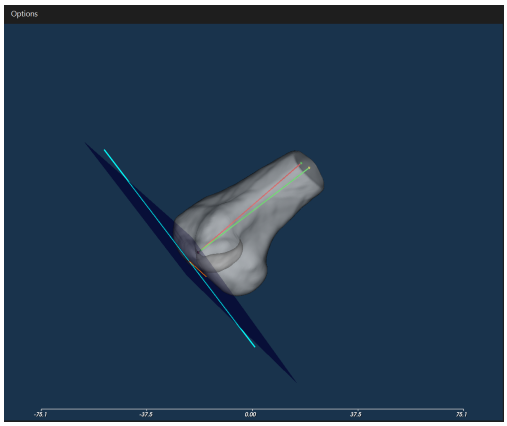


(b) Sidebar interface displayed during the cutting plane origin selection process

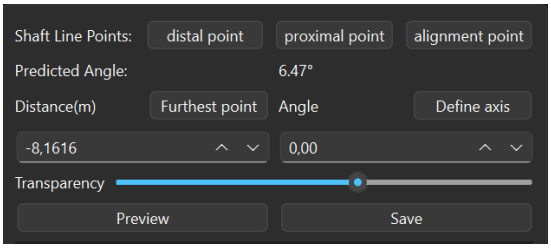
Figure 4.7: Definition of the cutting plane origin and corresponding side panel interface

To provide the user with dynamic control over the resection depth, the application supports two methods of translational alignment along the anatomical axis vector.

First, manual depth configuration can be executed via a dedicated distance spinner box interface in the side panel, which handles incremental linear translations of the plane along the axis. Second, to streamline the clinical workflow, the application incorporates an automated surface mesh evaluation tool. Upon triggering the furthest point function, the software executes a search across the vertices of the femoral STL surface mesh to identify the absolute outermost coordinate relative to the plane's current orientation. The cutting plane is then automatically translated to this anatomical boundary, establishing an accurate baseline for calculating the distal bone resection thickness, as shown in Figure 4.8.



(a) Selection of the furthest point for defining the cutting plane with simultaneous visualization of the virtual mechanical axis (green), the anatomical axis (red), the primary reference vector (orange), the rotation axis (blue) and the cutting plane (dark blue)



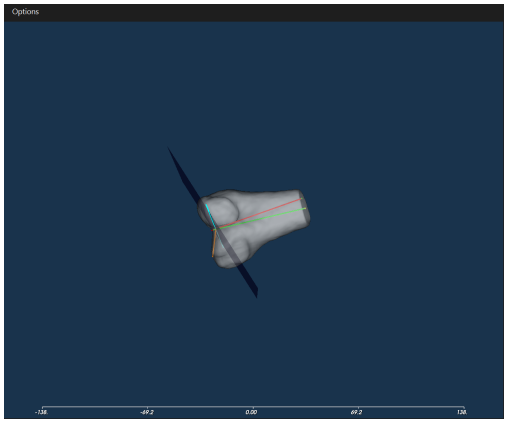
(b) Sidebar interface displayed during the furthest point selection process

Figure 4.8: Definition of the cutting plane using the furthest point and corresponding side panel interface

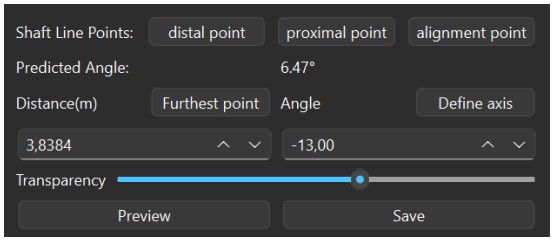
Following the translation of the resection plane to the anatomical furthest point, a localized coordinate system is constructed to facilitate controlled rotational adjustments. This system establishes two mutually orthogonal reference vectors within the resection plane:

The primary reference vector, represented as an orange axis in Figure 4.9a, is defined by connecting the intersection point of the anatomical axis and the resection plane directly to the designated anatomical furthest point. Subsequently, a secondary rotation axis, represented as a blue axis in the same figure, is computed. To ensure strict planarity, this rotation axis is defined as a line perpendicular to both the anatomical axis and the orange reference vector, effectively satisfying a cross-product relationship ($v_{\text{rotation}} = v_{\text{anatomical}} \times v_{\text{primary}}$). This kinematic constraint isolates the plane's rotational degrees of freedom, allowing the

user to execute angular adjustments without shifting the established distal resection depth as shown in Figures 4.9.



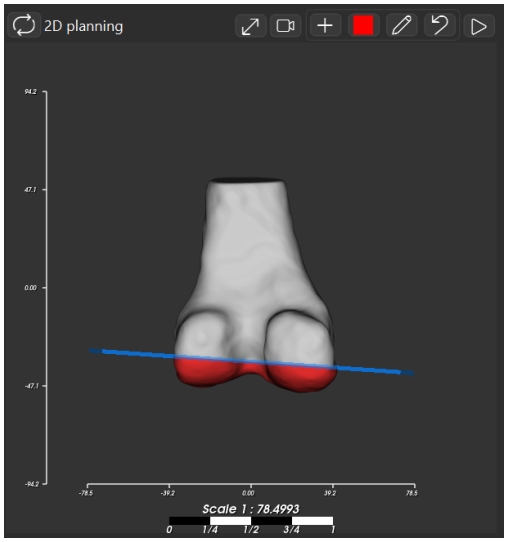
(a) Cutting plane after rotational adjustment during the surgical planning process (dark blue) with simultaneous visualization of the virtual mechanical axis (green), the anatomical axis (red), the primary reference vector (orange) and the rotation axis (blue)



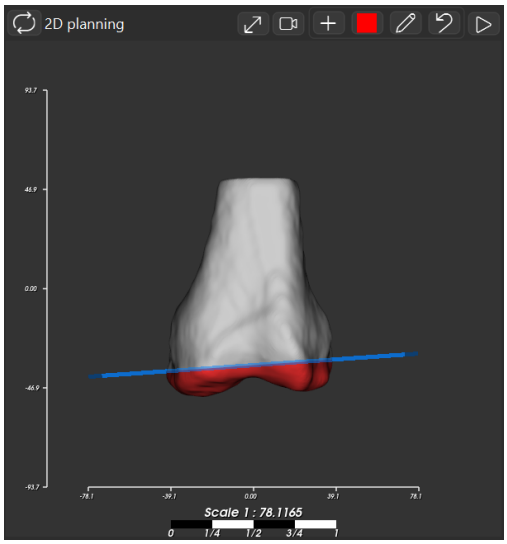
(b) Sidebar interface after defining the rotational adjustment parameters

Figure 4.9: Rotational adjustment of the cutting plane and corresponding side panel interface

In addition to constraining spatial rotations, the computed rotation axis serves as a geometric baseline for cross-sectional visualization. Two virtual cameras are positioned along this rotational axis to project the 3D anatomical model into synchronized 2D orthographic viewports that can be seen in the 2D Viewer. This dual-perspective setup generates inferior and superior 2D planar cross-sections of the distal femur. Within these 2D views, the bone volumes intersecting the planned resection plane are dynamically shaded in red to visually represent the projected bone resection zones, as shown in Figure 4.10.



(a) Inferior 2D view of the planned cutting plane



(b) Superior 2D view of the planned cutting plane

Figure 4.10: Inferior and superior 2D views of the planned cutting plane

To provide manual override capabilities if the initial coordinate system requires optimization, an axis redefine function is integrated. By activating this feature, the user can pick a new point within the 3D viewport in, with the primary reference vector being recalculated with the newly designated coordinate replacing the established anatomical furthest point in the vector definition.

To modify the interface layout during the planning stage a toggle switch allows the user to invert the positions of the 3D and 2D windows, maximizing the 2D layout when planar adjustments require a larger viewing area, as shown is Figure 4.11.

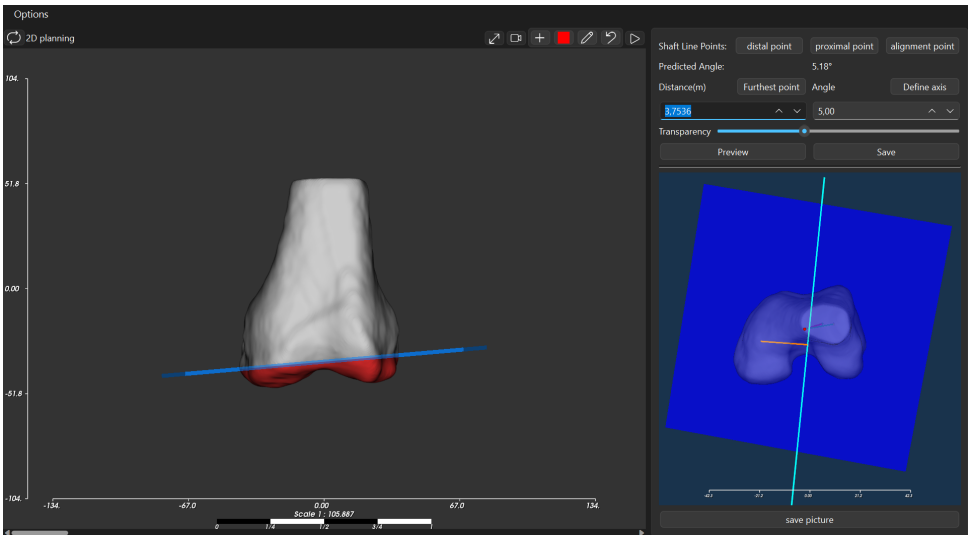


Figure 4.11: Workspace layout inversion maximizing the 2D planning window for detailed cross-sectional inspection

Furthermore, quantitative assessment utilities are integrated to enhance the 2D planning workflow. A digital ruler tool enables precise point-to-point distance measurements directly within the cross-sectional plane shown in Figure 4.12. Additionally, a 2D-to-3D reference annotation tool allows the user to plot discrete reference points within the 2D projection. The software then back-projects these coordinates into the 3D workspace via linear ray-casting, rendering corresponding spatial lines in the 3D view. This bi-directional visualization significantly improves the spatial correlation between the isolated 2D cutting planes and the global 3D anatomy shown in Figure 4.13.

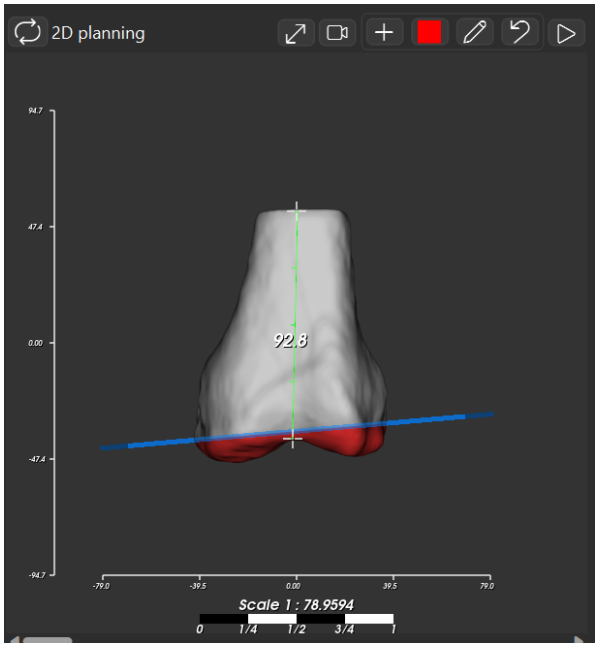


Figure 4.12: Distance measurement functionality available in the 2D planning view

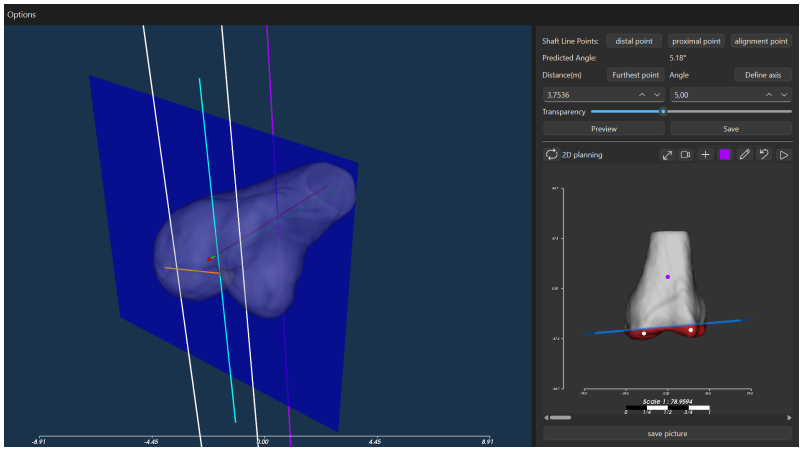


Figure 4.13: Reference point annotation feature in the 2D visualization environment

To translate the virtual surgical plan into a clinical execution strategy, the application incorporates a predictive instrumentation module. Activating the preview function dynamically

fits the 3D CAD model of the physical distal femoral cutting block onto the segmented femur mesh surface, as shown in Figure 4.14. This alignment pipeline uses sequential geometric constraints to position the instrument: First, the internal slot plane of the virtual cutting guide is constrained to be strictly coplanar with the user-defined resection plane. Second, a predefined central reference vector embedded within the guide's local coordinate system is aligned collinearly with the previously computed rotation axis.

To simulate physical seating onto the bone surface, a two-stage fitting algorithm is executed. First, the cutting guide is translated downward along the rotation axis until its internal surface registers a collision detection event at an initial contact point on the femoral STL mesh. Second, using a localized axis passing through this primary contact point and oriented normal to the cutting plane, the guide is rotated until a secondary surface location achieves contact with the bone mesh, establishing a stable, dual-point tangential seating orientation.

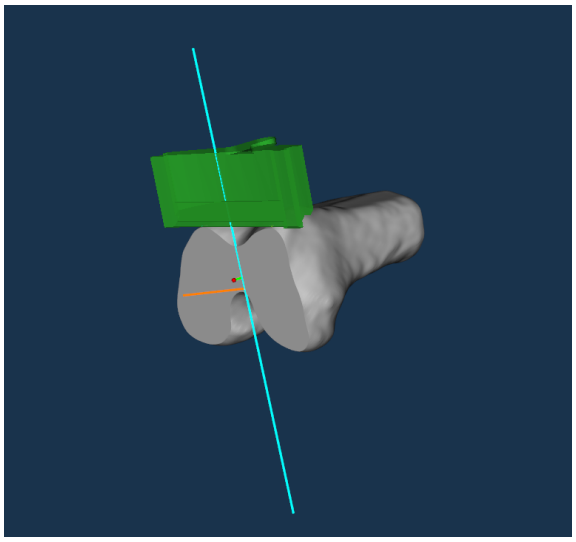


Figure 4.14: Virtual seating simulation of the localized femoral distal cutting block onto the bone mesh surface

Once the virtual fitting is complete, the save button allows the user to export the planning data. The software exports a post-resection femoral STL mesh reflecting the simulated bone cuts, the STL guide in the choosen position, and a plain text configuration file containing the 4×4 transformation matrix required to map the distal femoral cutting guide from its initial origin to its choosen position relative to the femur based on the surgical planning.

The physical cutting guide geometry utilized in this pipeline was modeled directly from a standard distal femoral instrumentation block provided by the *Centro de Simulação Biomédica do Hospital de Santo António*. Because the geometric constraints, alignment vectors, and internal seating tolerances within the pipeline are mathematically bound to this specific reference design, the generated spatial transformations are exclusive to this hardware configuration and cannot be generalized to alternative manufacturing layouts without retraining

the local coordinate origins.



(a) Distal femoral cutting guide front



(b) Distal femoral cutting guide back

Figure 4.15: Front and back views of the distal femoral cutting guide

4.5 Point Cloud Registration

The objective of the registration pipeline is to determine the rigid body transformation matrix, consisting of a rotation matrix R and a translation vector t, that aligns the pre-operative 3D femur model coordinate system with the real-time point cloud coordinate system acquired by the depth camera. This spatial alignment is achieved through a two-stage registration workflow: an initial coarse alignment followed by a fine surface registration that is continuously applied.

4.5.1 Data Acquisition System

To obtain the intraoperative spatial data, an Ensenso N36 depth camera system was utilized to provide dense 3D surface profiles. Prior to active tracking, the camera's intrinsic calibration and tracking baseline were validated across its working volume using the manufacturer's automated verification routine within the NxView calibration wizard. This evaluation analyzes a specialized Ensenso dot-grid calibration pattern (390×390 mm, 48 mm grid spacing) placed at varying depths to compute spatial and geometric discrepancies. The software interface during this verification process is shown in Figure 4.16.

The physical evaluation spanned distances ranging from 500 to 1000 mm relative to the sensor. Based on this pre-acquisition verification, the optimal workspace volume was constrained to a depth between 500 and 700 mm to minimize localized triangulation errors during the downstream procedures.

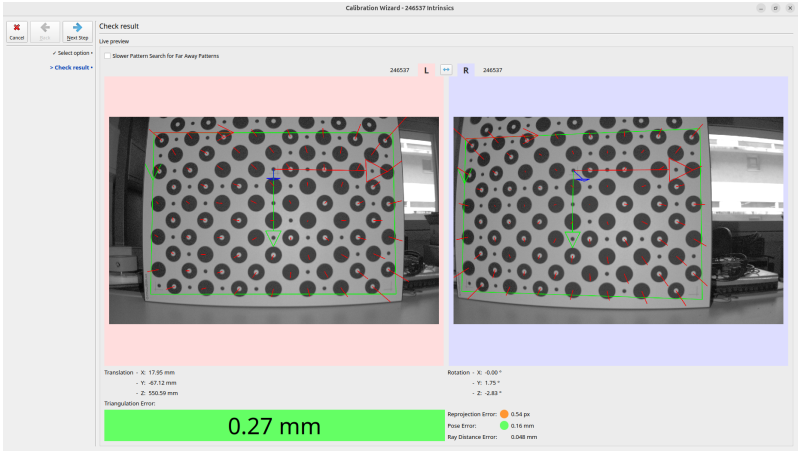


Figure 4.16: NxView calibration wizard interface during the camera validation routine

The sensor layer is integrated into a Robot Operating System 2 (ROS2) Jazzy framework via the official Ensenso driver. A dedicated rclpy node, designated CloudCheck, subscribes to two concurrent data streams, a `/point_cloud` topic, which provides raw 3D geometric coordinates of the scanned surface, and a `/rectified/left/image` topic, which supplies a rectified grayscale image capturing surface features.

An approximate time synchronization policy pairs these streams within a 0.05 s window to map the 3D points to their corresponding pixels. Synchronized frames are then filtered using spatial thresholds ($z \leq 0.75$ m and $y \geq -0.05$ m) to isolate the target anatomy. The remaining data is converted into an Open3D point cloud structure with mapped intensity values.

4.5.2 Coarse Alignment

Since local registration methods are highly sensitive to initialization and may converge to local minima when the initial pose estimate is poor, a coarse alignment stage is first performed to provide an approximate rigid transformation between the pre-operative model and the intraoperative point cloud [93].

The initialization procedure begins with the selection of anatomical landmarks on the pre-operative femur model using a dedicated landmark selection application. The application, developed using PyQt, VTK, and Open3D, provides an interactive environment in which the user loads the pre-operative point cloud and manually identifies anatomical landmarks, as shown in Figure 4.17.

To facilitate point selection, the input point cloud is first downsampled and converted into a continuous triangular surface mesh using Poisson surface reconstruction. Displaying a reconstructed surface instead of a sparse point cloud provides a smoother and denser representation of the anatomy, making it easier for the operator to accurately identify anatomical

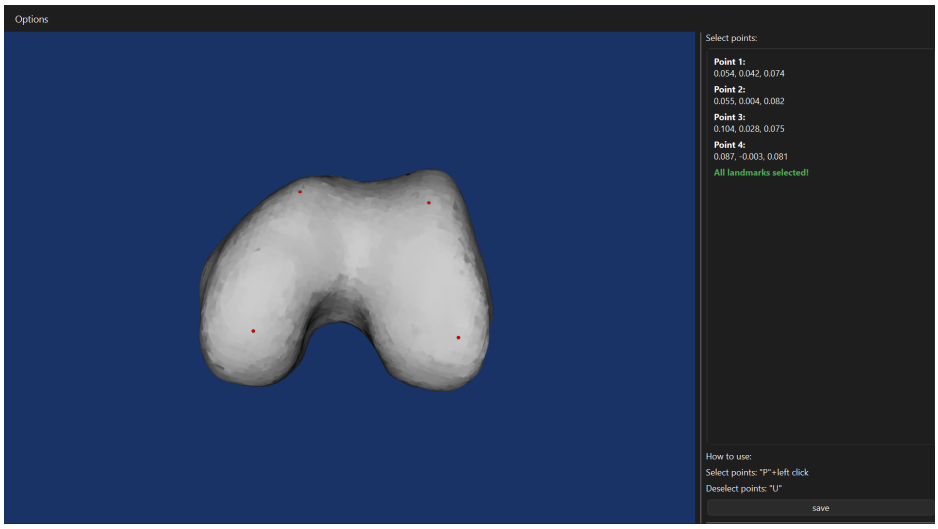


Figure 4.17: Pre-operative landmark selection application used to identify anatomical reference points on the femur model

features and select landmarks. The selected landmark coordinates are then exported to a text file and used as the source landmark set for the registration procedure.

During the intraoperative stage, a dedicated application is used to perform manual landmark matching and initialize the registration process. The application loads the pre-operative femur point cloud together with the previously selected landmark coordinates and acquires a live point cloud from the stereo camera through a ROS2 interface.

An Open3D based visualization tool then displays the pre-operative model and the acquired point cloud side by side, providing visual guidance for identifying corresponding anatomical features. The operator manually selects the landmarks on the target point cloud in the same sequential order as the pre-operative landmarks, as shown in Figure 4.18. A minimum of three corresponding landmarks is required to estimate the initial rigid transformation.

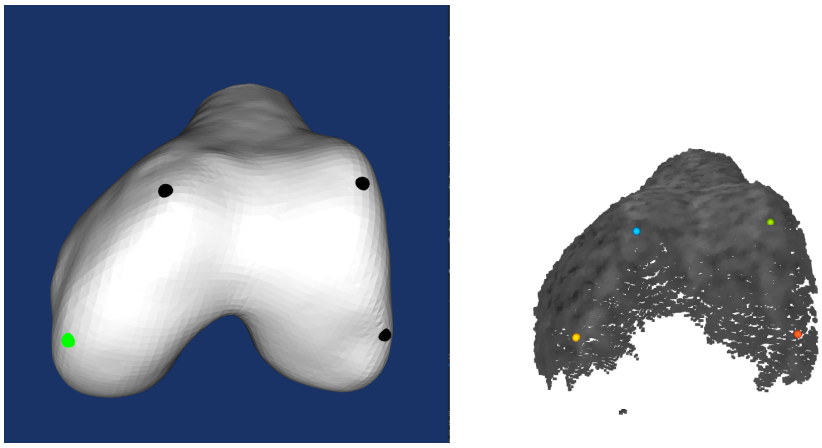


Figure 4.18: Intraoperative landmark matching interface showing the pre-operative femur model with selected landmarks (left) and the acquired point cloud (right)

To ensure responsive interaction and reduce computational cost, both point clouds are uniformly downsampled using a voxel grid size of 2.0 mm. Since the user-selected screen positions do not necessarily correspond to exact points in the downsampled cloud, a K-D-tree nearest-neighbour search is performed to associate each selection with its corresponding point index. The resulting set of landmark correspondences is then used by Open3D's `TransformationEstimationPointToPoint` method to compute an initial rigid transformation, which serves as the starting point for the subsequent fine registration stage. The resulting point clouds after the transformation obtained in the coarse alignment are shown in Figure 4.19

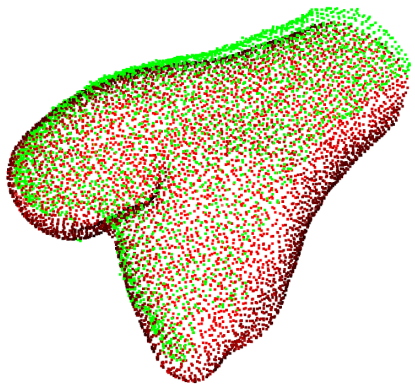


Figure 4.19: Point clouds after coarse alignment using the initial transformation estimate.

4.5.3 Fine Registration

Following the coarse alignment stage, a fine registration pipeline is applied to achieve the geometric accuracy required for surgical guide placement. This refinement is performed

using the point-to-plane ICP algorithm. Prior to optimization, both the source and target point clouds are downsampled using a voxel size of 0.25 mm to preserve fine anatomical detail while reducing computational cost and the intial transformation previously obtained from the coarse alignment is applied. Surface normals are then estimated using a hybrid K-D-tree neighborhood search with a radius of 1.0 mm and a maximum of 50 neighbouring points.

Point-to-plane ICP was selected instead of the classical Point-to-Point formulation because it generally exhibits faster convergence and improved accuracy when the initial alignment is already close to the optimum[43]. By minimizing the distance from each source point to the tangent plane defined by the corresponding target point normal, the algorithm better exploits the local surface geometry of the femur and is usually less sensitive to the anisotropic sampling density commonly observed in stereo-derived point clouds[43, 94].

At each iteration, correspondences are established by performing a nearest-neighbour search in the target point cloud. Correspondence pairs separated by more than 1.0 mm are rejected to prevent outliers and partial surface mismatches from influencing the optimization. The resulting rigid transformation is iteratively updated until either the convergence criteria are satisfied or the maximum number of iterations is reached.

Because the coarse alignment stage provides an initial transformation that already places the two point clouds in close proximity, the ICP optimization operates within the basin of attraction of the global minimum and reliably converges to the final registration matrix, denoted by T_{global}. The resulting point clouds after the transformation obtained in the ICP are shown in Figure 4.20.

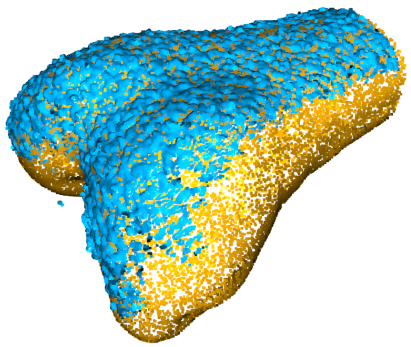


Figure 4.20: Point clouds after the ICP refinement

4.5.4 Real-Time Tracking

Following the computation of T_{global}, the registration process transitions into a real-time tracking mode. As new point clouds are continuously acquired from the stereo camera,

point-to-plane ICP is repeatedly executed to estimate the incremental rigid transformation between the current model pose and the newly acquired data. By performing registration between consecutive frames, only small inter-frame displacements need to be estimated, which improves the robustness of the optimization and helps maintain convergence to the correct alignment.

The resulting incremental transformation is used to continuously update the pose of the pre-operative model in real time, enabling the projected surgical guidance information to remain aligned with the patient throughout the procedure, as shown in Figure 4.21 .

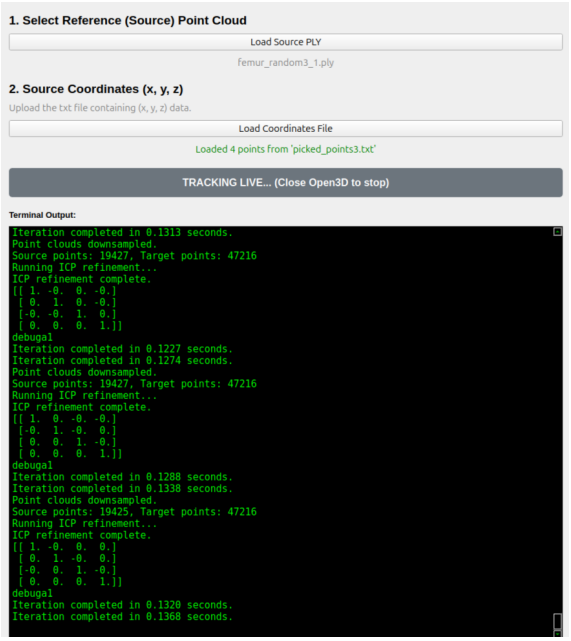


Figure 4.21: Graphical user interface used during real-time tracking, displaying the continuously the transformation matrix between the pre-operative model and the live point cloud

4.5.5 Validation

To comprehensively evaluate the accuracy, robustness, and clinical viability of the developed SAR registration framework, the validation was divided into three progressive experimental phases:

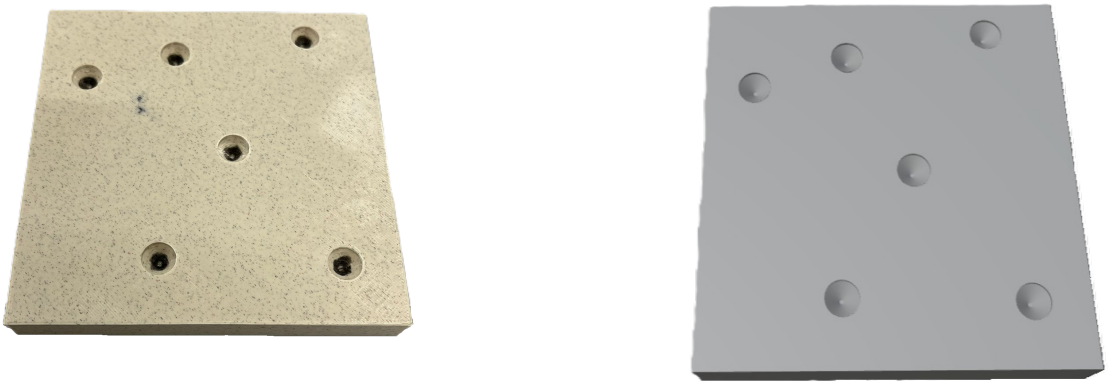
1. **Phase 1: Geometric Baseline Validation**, utilizing a custom calibration rig to isolate algorithmic execution from anatomical ambiguity.
2. **Phase 2: Morphological Surface Validation**, utilizing a knee phantom to evaluate inter-operator reproducibility and usability on anatomical contours.
3. **Phase 3: Intraoperative Target Validation**, assessing registration performance directly on bare bone morphology under both fully exposed and partially occluded surgical conditions.

Phase 1: Geometric Baseline Validation

To isolate the baseline accuracy of the registration pipeline from tracking errors or user landmark-selection uncertainties, a custom validation object was designed in CAD and fabricated via 3D printing.

The rig features a geometrically stable, non-symmetrical multi-planar structure embedded with 6 conical recesses acting as geometric targets with known ground-truth coordinates, as shown in Figure 4.22. To ensure an independent and unbiased evaluation, these landmarks were divided into two distinct groups:

- **Fiducial Points:** A subset of 3 landmarks distributed across the outer surfaces, utilized exclusively during the manual initialization phase to compute the initial coarse alignment and calculate the Fiducial Registration Error (FRE).
- **Target Points:** A separate subset of 3 hidden or internal landmarks, completely excluded from the registration optimization loops and reserved strictly for calculating the Target Registration Error (TRE).



(a) Photo of the validation object

(b) STL model of the validation object

Figure 4.22: Validation rig used in the study

Using the registration pipeline described in the previous sections, the registration accuracy was evaluated immediately after the coarse alignment stage and after the initial ICP refinement, before entering the real-time tracking mode. The performance of the pipeline was quantified using four complementary metrics: FRE, TRE, inlier root mean square error (RMSE), and fitness score. The FRE measures the residual distance between the corresponding fiducial landmarks used during registration, whereas the TRE measures the error at independent target landmarks that were not involved in the registration process, providing an unbiased estimate of the registration accuracy. The inlier RMSE represents the root mean square distance between all corresponding point pairs classified as inliers, while the fitness

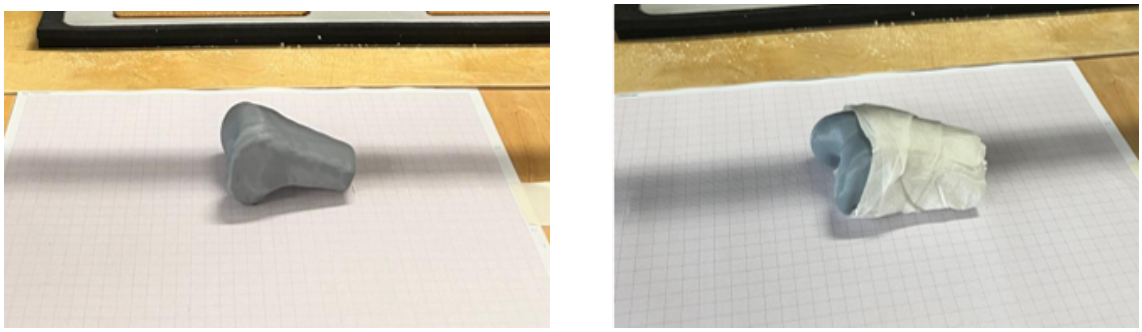
score corresponds to the proportion of source points that find a valid correspondence within a maximum distance threshold.

Phase 2: Morphological Surface Validation

A knee phantom was used to evaluate the usability and repeatability of the proposed registration pipeline. Multiple users independently performed the complete registration workflow, including the manual landmark selection procedure described. The resulting registrations were evaluated using the same performance metrics (FRE, TRE, inlier RMSE, and fitness score) to assess the influence of user variability on the registration outcome.

Phase 3: Intraoperative Target Validation

A final validation was performed using two independent sets of three 3D-printed femurs. The corresponding pre-operative anatomical femoral models were derived from volumetric CT scans. Following semi-automated segmentation in 3D Slicer, the extracted femoral geometries were exported as STL meshes. To interface with the registration pipeline, each mesh was uniformly scaled relative to the global origin and uniformly sampled using Open3D to generate dense 3D reference point clouds in polygon file format. In the first set, the complete femur surface was exposed, resulting in a registration scenario where the pre-operative model closely matched the visible anatomy. In the second set, the exact same three femurs were reused but covered with paper drapes around the shaft (designated with an "m" suffix), leaving only the more distal region exposed, as shown in Figure 4.23. This configuration was designed to emulate the clinical conditions encountered during TKA, where only a limited portion of the distal femur is visible after surgical exposure, resulting in a partial-to-partial registration problem. For both datasets, the registration quality was evaluated using the inlier RMSE and fitness metrics with a maximum correspondence distance of 0.5 mm. In addition, a qualitative visual assessment was performed to verify the anatomical alignment between the registered pre-operative model and the acquired point cloud.



(a) Bone with complete surface exposure

(b) Bone with partial surface exposure due to paper drapes (m)

Figure 4.23: Experimental surface visibility conditions of the 3D-printed femurs used for registration validation

4.6 Projector-Camera Calibration

To accurately project surgical guide boundaries onto the physical environment, the intrinsic parameters and extrinsic geometric relationship between the depth sensor and the projector must be established.

4.6.1 Intrinsic Parameters and the Pinhole Model

The imaging geometry of a standard camera is modeled using the idealized pinhole camera projection. Under this model, a 3D point in the camera coordinate frame, denoted as $P_c = [X_c, Y_c, Z_c]^T$, is projected onto the 2D image plane at pixel coordinates $p = [u, v]^T$. This transformation is expressed in homogeneous coordinates as:

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = K \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} \tag{4.1}$$

where $s = Z_c$ is the projective scale factor representing the depth along the optical axis, and K is the camera intrinsic parameter matrix, assuming zero skew, defined as:

$$K = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \tag{4.2}$$

In this matrix, f_x and f_y denote the camera focal lengths expressed in pixel units along the horizontal and vertical axes, respectively, while c_x and c_y represent the pixel coordinates of the principal point, the orthogonal projection of the camera optical center onto the image plane.

4.6.2 Planar Calibration and Zhang's Methodology

The extraction of the camera intrinsic parameters relies on the planar calibration methodology proposed by Zhang [66]. This approach eliminates the need for expensive 3D calibration apparatus by employing a planar target with a known geometric pattern, such as a checkerboard.

Assuming without loss of generality that the planar target lies on the $Z = 0$ plane in the world coordinate system, the mapping between planar points and image coordinates reduces to a 2D homography H. By capturing the planar target across multiple distinct spatial orientations ($N \geq 3$), each view provides two independent algebraic constraints on the intrinsic parameters derived from the orthonormality of the rotation matrix columns. This system of linear equations provides an initial closed-form estimate of K, which is subsequently refined

alongside extrinsic poses through non-linear Levenberg-Marquardt optimization minimizing the reprojection error across all observed corners.

4.6.3 Lens Distortion and Image Rectification

Physical camera lenses deviate from the idealized pinhole model by introducing geometric distortions that displace image pixels from their true perspective coordinates. According to the Brown–Conrady model, this is categorized into:

- **Radial distortion** (k_1, k_2, k_3)**:** Caused by symmetric lens curvature imperfections, resulting in barrel or pincushion warping away from the optical center.
- **Tangential distortion** (p_1, p_2)**:** Caused by physical decentering or misalignment of lens elements relative to the image sensor plane.

In medical navigation and surgical tracking pipelines, uncorrected optical distortion introduces systematic spatial biases into point cloud reconstruction and coordinate registration. To eliminate these errors, the estimated distortion coefficients are applied in an inverse coordinate mapping pipeline. For each pixel coordinate in the target undistorted frame, the corresponding distorted location on the raw sensor grid is computed and sampled via bilinear interpolation. This rectification ensures straight lines in the physical workspace remain straight in the digital image, preserving the linear assumptions of the pinhole model throughout downstream spatial tracking.

4.6.4 Ensenso Depth Camera Intrinsic Parameters

For the Ensenso N36 camera, the planar calibration methodology described in the previous section was bypassed in favor of the factory-calibrated parameters provided by the manufacturer. Rather than executing an independent camera calibration procedure, the factory-rectified intrinsic parameters and distortion coefficients of the left camera sensor were extracted directly from the device firmware via the Ensenso NxTree.

Because the sensor assembly's native stereo matching and depth map generation are co-registered with the optical center of the rectified left camera, the resulting 3D point cloud and the 2D monochrome imagery share a unified spatial reference frame. And since the internal stereo pipeline rectifies lens aberrations before outputting images, the effective lens distortion vector is zero. Here, K_{cam} is the camera intrinsic matrix, while d_{cam} represents the lens distortion coefficients:

$$K_{\text{cam}} = \begin{bmatrix} 2256.487 & 0 & 872.692 \\ 0 & 2256.487 & 591.342 \\ 0 & 0 & 1 \end{bmatrix}, \quad d_{\text{cam}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \end{bmatrix}^T \tag{4.3}$$

4.6.5 Projector Calibration

The projector is modeled as an inverse pinhole camera. The estimation of the projector's intrinsic parameters and its extrinsic spatial orientation relative to the camera (R_{cam}^{proj}, t_{cam}^{proj}) was performed based on the methodology described by Moreno and Taubin [67].

The calibration pipeline was structured as follows:

Chessboard pattern: A planar checkerboard calibration target comprising 12 × 9 squares, each with a side length of 15 mm, was generated using the Calib.io calibration pattern generator. The pattern was printed at its intended scale and mounted onto a rigid cardboard backing to ensure planarity and prevent deformation during image acquisition. The resulting calibration target is shown in Figure 4.24.

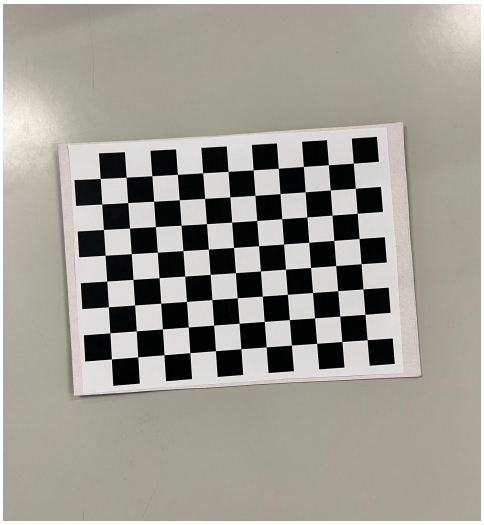


Figure 4.24: Chess board pattern

Structured Light Pattern Projection: Based on the projector resolution, a sequence of 40 orthogonal Gray-code structured light patterns was generated using `structured_light` module (`GrayCodePattern`) from OpenCV. In addition, two reference images corresponding to uniform minimum (black) and maximum (white) projector illumination were generated to establish the lower and upper illumination bounds for pattern decoding for a projector resolution of 1024 x 768 and projected onto a static, planar checkerboard target in 10 different positions: one flat on the table, 4 inclined in each direction in the central position of the camera view, 4 more inclined closer to the sides of the camera view and one flat closer to the camera. The left camera of the Ensenso N36 captures the images of the checkerboard under these projected patterns for each projection and pattern position, as shown in Figure 4.25.

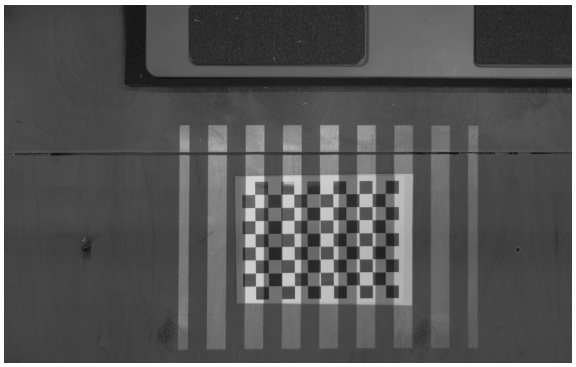


Figure 4.25: Projected Gray-code structured light pattern captured by the left camera

To achieve a valid calibration result, the projector and camera settings were modified to ensure the printed checkerboard corners and the projected Gray code patterns were distinct in the captured images. The projector brightness was reduced because high illumination levels caused specular reflections on the target surface, resulting in overexposed regions where the printed checkerboard pattern could not be detected. The camera exposure time was synchronized with the projector refresh rate to prevent flickering between pictures and rolling shutter artifacts, which distort the captured binary patterns.

Pixel-to-Pixel Correspondence Decoding and Calibration Execution: To establish the mapping between the camera and projector frames, a shadow mask is first generated by computing the absolute difference between images captured under full illumination and zero illumination. A binary threshold of 30 is applied to this differential image to isolate pixels receiving direct projector illumination from ambient or shadowed regions.

The checkerboard corners are detected within the illuminated workspace using the OpenCV's `findChessboardCorners` function target grid configuration of 11×8 corners. To achieve sub-pixel accuracy, the detected corner locations are refined using the `cornerSubPix` function with an 11×11 pixel search window.

For each refined sub-pixel camera corner that falls within the valid illumination mask, the underlying sequence of 40 orthogonal Gray code patterns is evaluated. Because a camera corner rarely projects directly onto an exact integer projector pixel coordinate, and to recover from local decoding failures caused by shadows or surface transitions, a local spatial mapping is executed.

A 7×7 pixel neighborhood (corresponding to a window radius of 3 pixels) is evaluated around each corner. Every pixel within this window that successfully decodes a Gray code sequence yields a point correspondence pair. If a minimum of 12 valid neighbor points are collected, a local homography matrix H is computed between the camera and projector coordinate subsets using a Random Sample Consensus algorithm with a reprojection error threshold of 0.5 pixels to eliminate noise and decoding outliers. The resulting homography

is applied via a perspective transformation to map the sub-pixel camera corner coordinate precisely into the 1024×768 projector display plane.

Once the full set of point correspondences is populated across all captured target poses, calibration is executed in a two-stage pipeline:

1. **Projector Intrinsic Estimation:** The 3D world coordinates of the target corners (defined based on a known square size of 15.0 mm) are mapped to their calculated projector pixel coordinates. The projector's intrinsic matrix K_{proj} and distortion coefficients are estimated using the `calibrateCamera` function, initialized with a principal point guess based on the projector resolution.
2. **Stereo System Calibration:** The spatial relationship between the devices is computed using the `stereoCalibrate` function. The intrinsic parameters of the left Ensenso camera sensor are held fixed to preserve the factory alignment with the depth reference framework. The optimization yields the final rotation matrix R and translation vector T, defining the transformation between the camera origin (O_{cam}) and the projector.

The performance of the projector-camera calibration pipeline was evaluated through quantitative numerical metrics and qualitative closed-loop projection tests. The stereo calibration algorithm processed structured light acquisitions by fixing the Ensenso camera's factory intrinsics (`CALIB_FIX_INTRINSIC`) while optimizing the projector's internal parameters and relative spatial extrinsics. Radial distortion coefficients (k_1, k_2, k_3) were fixed to zero during optimization (`CALIB_FIX_K1`, `CALIB_FIX_K2`, `CALIB_FIX_K3`), while tangential distortion parameters (p_1, p_2) were optimized. The reprojection error (global stereo RMSE) was calculated to assess calibration accuracy.

To validate the spatial accuracy of the estimated parameters, a closed-loop back-projection experiment was performed. The checkerboard target was localized in 3D camera coordinates by solving the Perspective-n-Point problem using corner detections from OpenCV's `findChessboardCorners`. Using the optimized extrinsic (R, T) and intrinsic $(K_{\text{proj}}, D_{\text{proj}})$ matrices, predicted corner positions were mapped to the projector frame to generate a synthetic canvas of green circular markers with a 4-pixel radius. This synthetic pattern was subsequently transmitted to the projector, cast back onto the physical checkerboard, and visually evaluated.

4.7 Spatial Augmented Navigation System Implementation

4.7.1 Hardware Setup and Configuration

To evaluate the spatial augmented reality system under conditions that simulate a clinical environment, a dedicated hardware configuration was developed. In an operating room, the overhead space directly above the surgical site is restricted due to the presence of surgical illumination systems, as shown in Figure 4.26. Consequently, the navigation system must operate at an oblique angle relative to the anatomical target.



Figure 4.26: Operating room configuration and equipment layout during conventional TKA [6]

To replicate these spatial constraints, the projector and the Ensenso N36 depth camera were mounted side-by-side on a rigid horizontal fixture. The target surface of the femur was positioned below the sensor assembly and angled relative to the instruments. This orientation simulates the line-of-sight and projection angles required to operate within the physical limitations of a standard surgical workspace, as shown in Figure 4.27.

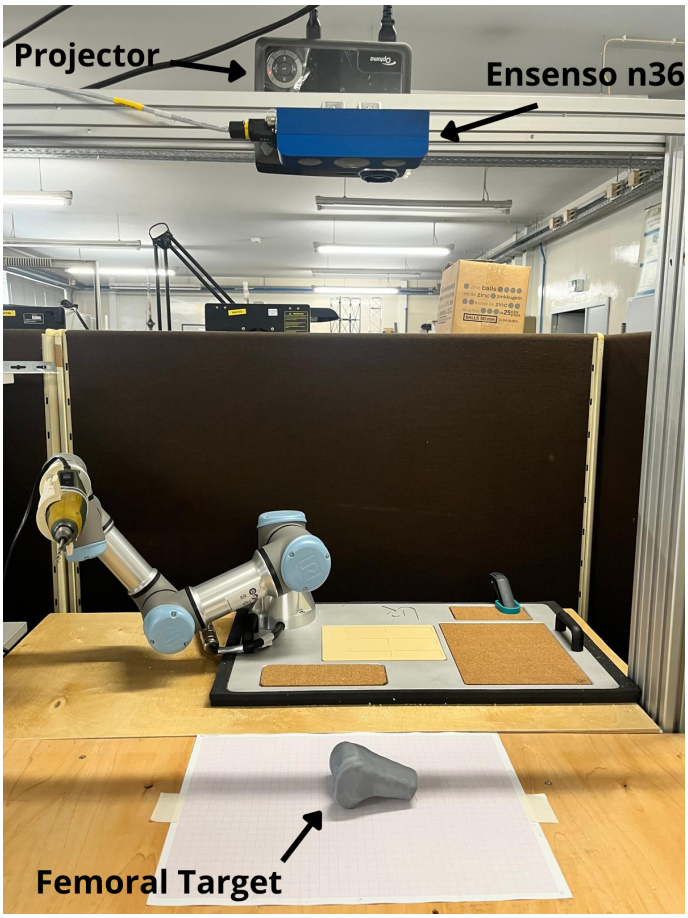


Figure 4.27: Experimental hardware setup for the SAR system, showing the side-by-side mounting of the projector and Ensenso N36 depth camera above the femoral target

4.7.2 Surgical Guide Support Configuration

To prevent displacement of the surgical guide during the resection phase and to maintain geometric accuracy, a mechanical support assembly was implemented. The configuration consists of an adjustable articulating mechanical arm attached to a custom 3D printed adapter developed in SolidWorks. This adapter interfaces directly with the distal end of the mechanical arm and secures the surgical guide. This assembly maintains the spatial position of the instrument without manual assistance while it is rigidly fixed to the bone, as shown in Figure 4.28.



Figure 4.28: Mechanical support assembly for surgical guide stabilization, consisting of an articulating mechanical arm and a 3D-printed adapter

4.7.3 Intraoperative Alignment, Tracking, and Projection Pipeline

To overlay the planned surgical guide boundaries accurately onto the physical target during TKA, a comprehensive application (referred to as the Projection Application) was developed to execute a dynamic coordinate transformation and projection pipeline. This pipeline bridges the virtual planning environment with the physical workspace by synchronizing three distinct coordinate frameworks: the pre-operative planning space, the intraoperative camera space, and the projector display coordinate plane.

The spatial arrangement of the surgical instrument relative to the patient's anatomy is initially established within the pre-operative surgical planning application. The system imports this plan as a rigid-body homogenous transformation matrix (T_{plan}), alongside the point clouds of the outer border (P_{outer}) and inner border (P_{inner}) of the surgical guide's upper surface, as shown in Figure 4.29.

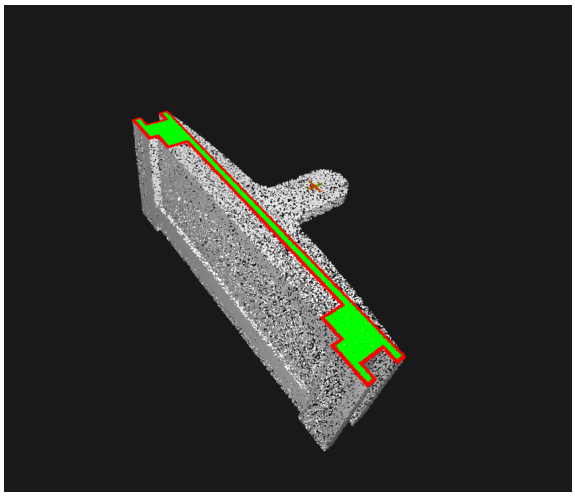


Figure 4.29: Boundary point clouds overlaid on the distal femoral cutting block model, showing the outer boundary (P_{outer}) highlighted in red and the inner boundary (P_{inner}) highlighted in green

Before real-time tracking begins, this pre-operative planning matrix is applied to translate and rotate the virtual guide borders into structural alignment with the reference bone point cloud:

$$P'_{tool} = T_{plan} \cdot P_{tool} \tag{4.4}$$

where $P_{tool} = [P_{inner}, P_{outer}]^T$. Because structural geometries exported from surgical CAD and planning environments typically utilize millimeters as their base unit, an isotropic scaling matrix is applied around the coordinate origin using a down-scaling factor of 0.001. This converts the point cloud coordinates into meters, enforcing unit consistency with the spatial metric domain of the Ensenso depth camera.

4.7.3.1 Dual-Stage Intraoperative Registration and Real-Time Tracking

The intraoperative localization pipeline relies on a continuous 3D point cloud stream from the Ensenso N36 depth camera. Spatial initialization follows the dual-stage registration methodology detailed in Section 4.5, where landmark-based manual alignment provides a coarse seed that is subsequently refined using point-to-plane ICP to establish the global base transformation T_{global} that is applied to the inner and outer guide border point cloud.

Once initialized, the system transitions into real-time tracking. For each incoming point cloud frame k, point-to-plane ICP registers the current guide state against the newly acquired point cloud to compute an incremental rigid-body delta transform. This running transformation matrix is applied in real time to the scaled surgical guide border geometries (P'_{inner}, P'_{outer}), tracking the physical bone's position dynamically as it shifts within the operational workspace.

$$P_{guide}^{(k)} = \Delta T_{ICP}^{(k)} P_{guide}^{(k-1)} \tag{4.5}$$

To prevent tracking divergence caused by hands or surgical instruments obstructing the line of sight, an integrated occlusion guardrail monitors the incoming point cloud and instantly pauses the continuous ICP loop locking the system into its last valid state $T_{\text{track}}^{(k-1)}$, whenever a spatial cluster exceeding 50 points violates an absolute depth threshold of 0.6 m.

4.7.3.2 Three-dimensional to Two-dimensional Perspective Projection Mapping

Once the tools' 3D positions are transformed into the active camera frame, the final step maps these coordinates to the projector's spatial hardware plane to create the augmented reality overlay. The projector is mathematically treated as an inverse pinhole camera system.

Let $X_c = (X_c, Y_c, Z_c)^T$ represent a 3D point from the tracked surgical guide boundaries in the camera frame. This point is mapped to the projector plane by applying the rigid-body extrinsic rotation matrix R_{cam}^{proj} and translation vector t_{cam}^{proj} obtained during calibration:

$$X_p = R_{cam}^{proj} X_c + t_{cam}^{proj} \tag{4.6}$$

The non-linear lens distortion profile of the projector is corrected using a five-element radial and tangential distortion vector $D = [k_1, k_2, p_1, p_2, k_3]$. The undistorted normalized projective coordinates (x', y') are defined as:

$$x' = \frac{X_p}{Z_p}, \quad y' = \frac{Y_p}{Z_p} \tag{4.7}$$

To correct for optical aberrations introduced by the projector lens, the normalized coordinates are adjusted using a five-parameter distortion vector $D = [k_1, k_2, p_1, p_2, k_3]$, representing radial (k_1, k_2, k_3) and tangential (p_1, p_2) distortion coefficients. With the radial distance defined as $r^2 = x'^2 + y'^2$, the distorted normalized coordinates (x'', y'') are computed via the Brown–Conrady model:

$$\begin{aligned} x'' &= x' \left(1 + k_1 r^2 + k_2 r^4 + k_3 r^6\right) + 2p_1 x' y' + p_2 \left(r^2 + 2x'^2\right) \\ y'' &= y' \left(1 + k_1 r^2 + k_2 r^4 + k_3 r^6\right) + p_1 \left(r^2 + 2y'^2\right) + 2p_2 x' y' \end{aligned} \tag{4.8}$$

Finally, the distorted normalized coordinates are mapped to pixel coordinates (u_p, v_p) on the 1024×768 display array using the projector's intrinsic calibration matrix K_{proj}:

$$\begin{pmatrix} u_p \\ v_p \\ 1 \end{pmatrix} = K_{proj} \begin{pmatrix} x'' \\ y'' \\ 1 \end{pmatrix} = \begin{pmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x'' \\ y'' \\ 1 \end{pmatrix} \tag{4.9}$$

where f_x, f_y denote the focal lengths and (c_x, c_y) is the principal point.

Using the OpenCV function `projectPoints`, the coordinates are converted into precise pixel array indices. The system renders the tool path into a custom multi-channel graphic buffer array: points representing the inner boundary (P_{inner}) are color-coded green, while points from the outer guide boundary (P_{outer}) are colored red to delineate tool margins. This coordinate buffer is mapped directly to a fullscreen hardware window projected down onto the target, as shown in Figure 4.30, allowing the surgeon to position the distal femoral cutting guide in the correct position choosen in the pre-operative planning application, as shown in Figure 4.31.

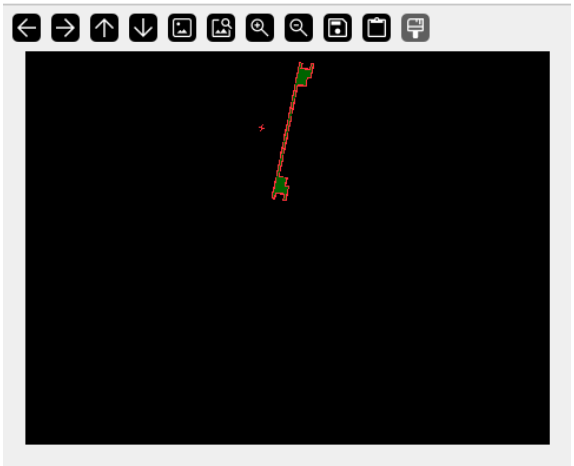
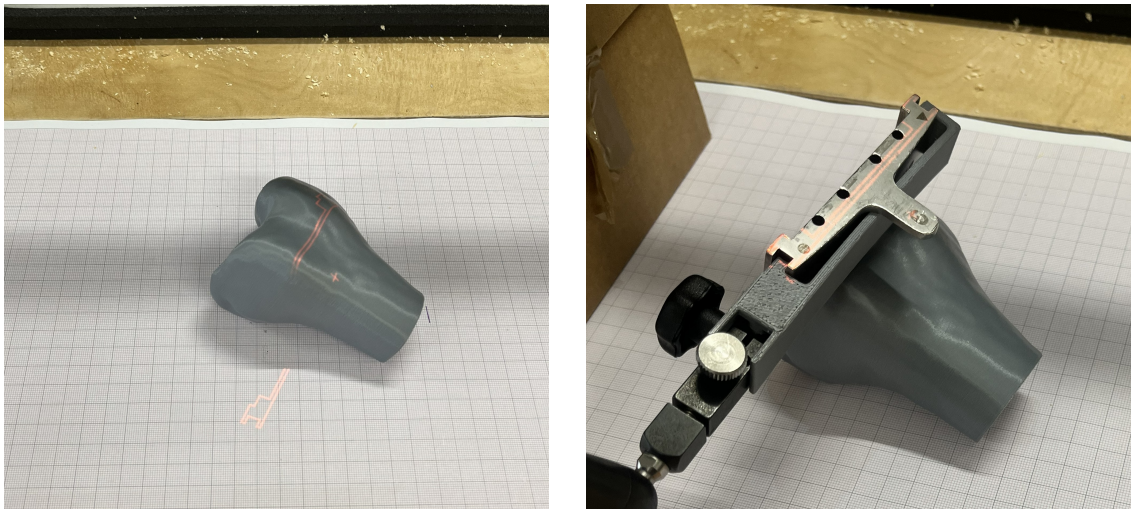


Figure 4.30: Window showing the 2D image of the surgical guide borders being projected



(a) Projection without guide.

(b) Projection with guide.

Figure 4.31: Comparison of the projection with and without the guide.

4.7.4 Pipeline Validation Methodology

A validation protocol was established to evaluate the accuracy of the projection and positioning pipeline by comparing the geometric relationships from the virtual planning stage against physical bench measurements. The three 3D-printed femurs used in this experiment were the same models previously utilized in the registration validation phase, ensuring anatomical consistency across testing stages.

To establish the reference points in the planning environment, three non-collinear points (F_1, F_2, F_3) were selected on the femur STL model in a VTK interface with `vtkCellPicker` set to a tolerance of 0.0005. Each selected position was visually verified by rendering a 1.0 mm red sphere at the picked site. Similarly, four points corresponding to the extremities of the surgical guide were extracted, as shown in Figure 4.32.

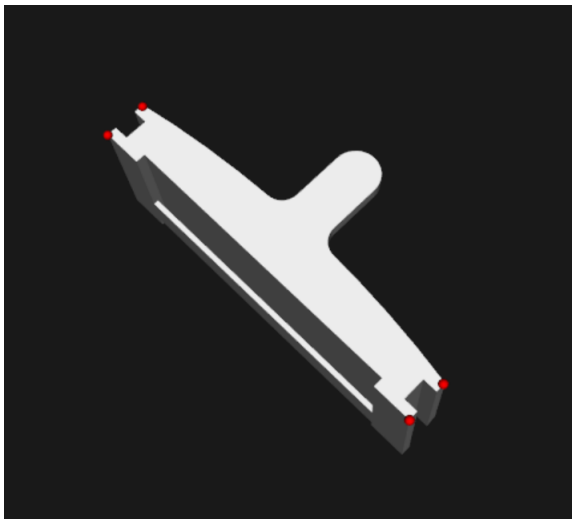


Figure 4.32: Selected guide points in the distal femoral guide 3D model

The transformation matrix obtained from the surgical planning application was then applied to these guide points to place the virtual surgical guide in its planned operational position relative to the femur, as shown in Figure 4.33. These expected spatial relationships served as the ground truth.

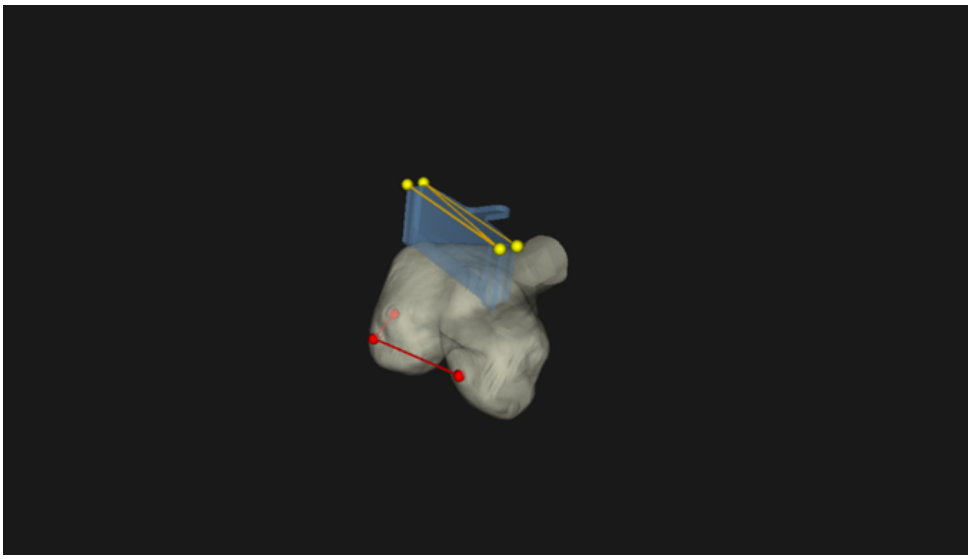


Figure 4.33: 3D models of the femur and surgical guide in the planned operational position, showing the selected femoral (red) and guide (yellow) points used to define their expected spatial relationship

For the physical setup, the three selected femur points were mapped and marked with a red marker on the 3D-printed bone to guarantee exact point correspondence during manual measurement. For the surgical guide, manual marking was unnecessary because the chosen extremities consisted of easily distinguishable physical boundaries. The surgical guide was then projected and physically positioned over the 3D-printed femur according to the system pipeline.

Physical measurements of the marked points were collected using a localized setup on the workbench. The horizontal X and Y values were read from a millimetric paper grid fixed to the table surface directly beneath the femur model, while the vertical Z values were captured using an orthogonal ruler.

Geometric features were calculated by determining two relative geometric metrics for both the expected planning configuration and the physical measurements: the spatial Euclidean distances between each femur-guide point pair, and the angle between the plane established by the three femur points and the directional vector of each femur-guide pair. Because these metrics rely on internal geometric relationships rather than an absolute global reference frame, they remain invariant to the exact placement of the models on the table.

Mathematically, the ground truth and physical benchtop geometric metrics were computed as follows. Let $F_i = [X_{F_i}, Y_{F_i}, Z_{F_i}]^T$ (for $i \in \{1,2,3\}$) represent the coordinates of the reference femur points, and let $G_j = [X_{G_j}, Y_{G_j}, Z_{G_j}]^T$ (for $j \in \{1,2,3,4\}$) represent the surgical guide extremity points. For each point pair (F_i, G_j), the 3D Euclidean distance d_{ij} is defined as:

$$d_{ij} = \|G_j - F_i\|_2 = \sqrt{(X_{G_j} - X_{F_i})^2 + (Y_{G_j} - Y_{F_i})^2 + (Z_{G_j} - Z_{F_i})^2} \tag{4.10}$$

The reference plane established by the three non-collinear femur points has a normal vector $n_F = (F_2 - F_1) \times (F_3 - F_1)$. The inclination angle θ_{ij} between the directional vector $v_{ij} = G_j - F_i$ and the femoral reference plane is calculated using the vector dot product:

$$\theta_{ij} = \arcsin\left(\frac{|v_{ij} \cdot n_F|}{\|v_{ij}\|_2 \|n_F\|_2}\right) \times \frac{180^\circ}{\pi} \tag{4.11}$$

For a given metric error across $N = 12$ evaluations per specimen (comprising 3×4 point pair combinations), the system's directional bias is quantified by the Signed Mean Error (SME), and overall magnitude error is quantified by the Mean Absolute Error (MAE):

$$\text{SME} = \frac{1}{N}\sum_{k=1}^{N} \Delta_k, \qquad \text{MAE} = \frac{1}{N}\sum_{k=1}^{N} |\Delta_k| \tag{4.12}$$

In addition to the invariant relative metrics, an absolute spatial validation was conducted to directly evaluate positioning accuracy by registering the virtual planning coordinates into the physical benchtop reference frame. A rigid point-based registration using singular value decomposition was applied to compute the optimal rotation and translation between the planned femur reference points and the measured benchtop points. This transformation was then applied to map the planned surgical guide points directly into the physical measurement space. The positioning error was evaluated by calculating the three-dimensional Euclidean distance between each transformed planned guide point and its corresponding physically measured position. Finally, the overall spatial accuracy across the guide boundaries was characterized by determining the average and maximum displacement errors.

Chapter 5

Results and Discussion

5.1 Camera Error Analysis

To validate tracking accuracy, the Ensenso N36 stereo camera error was evaluated across its working volume. The evaluation utilizes the camera's automated verification routine via the NxView calibration wizard, which analyzes a specialized Ensenso dot-grid pattern at varying depths to compute geometric and spatial discrepancies.

The tracking validation results from 500 mm to 1000 mm are summarized in Table 5.1.

Table 5.1: Tracking validation results

Distance (mm)	Triangulation error (mm)	Pose error (mm)	Ray distance error (mm)	Reprojection error (px)
500	0.24	0.15	0.053	0.60
550	0.26	0.18	0.047	0.54
600	0.24	0.16	0.049	0.58
650	0.24	0.16	0.047	0.59
700	**0.27**	**0.15**	**0.075**	**0.64**
750	0.29	0.17	0.077	0.66
800	0.32	0.18	0.070	0.68
850	0.34	0.20	0.066	0.69
900	0.38	0.22	0.077	0.71
1000	0.44	0.28	0.084	0.74

The data shows that errors scale proportionally with distance, with the triangulation error rising from 0.24 mm to 0.44 mm. Based on these metrics, a nominal working distance of approximately 700 mm was selected for the current calibration. At this distance, the system optimizes the tracking field of view while maintaining a low triangulation error of 0.27 mm and a pose error of 0.15 mm. Furthermore, this distance establishes an adequate

clearance envelope above the patient's knee, ensuring the camera hardware does not obstruct the surgical field and providing the surgeon with sufficient leverage during surgery.

5.2 Registration Pipeline Validation

To comprehensively evaluate the accuracy, robustness, and clinical viability of the developed SAR registration framework, three progressive validation phases were executed

5.2.1 Phase 1: Controlled System Validation

The quantitative metrics evaluated during the controlled validation phase are compiled in Table 5.2, contrasting the tracking behavior of the initial landmark-based coarse alignment against the final automated point-to-plane ICP refinement.

Table 5.2: Quantitative registration metrics before and after point-to-plane ICP Refinement

Registration Metric	Coarse Alignment	ICP-Refined Alignment	Points Used
FRE (mm)	1.356 ± 0.338	1.930 ± 0.312	3
TRE (mm)	1.932 ± 0.552	1.938 ± 0.532	3
Fitness (%)	100.0 ± 0.000	100.0 ± 0.000	> 1000
Inlier RMSE (mm)	1.764 ± 0.047	0.700 ± 0.024	> 1000

During initial manual registration, the workflow established viable baseline precision, holding target displacement under 2 mm (TRE of 1.932 ± 0.552 mm; FRE of 1.356 ± 0.338 mm) with complete ($100.0 \pm 0.000\%$) surface fitness across coarse and iterative stages.

The following ICP execution drove down the inlier RMSE from 1.764 ± 0.047 mm to 0.700 ± 0.024 mm, yet both FRE (1.930 ± 0.312 mm) and TRE (1.938 ± 0.532 mm) trended slightly upward. Rather than signaling poorer spatial alignment, this discrepancy reflects fundamentally competing mathematical loss functions between stages. The coarse step computes an analytical rigid transformation anchored directly to discrete landmark pairs, forcing spatial convergence at those specific sites and absorbing manual point-picking errors within a localized fit. In contrast, ICP replaces these discrete fiducial constraints with a continuous global distance minimization across the broader point cloud, prioritizing overall surface topology.

This shift toward a global geometric minimum unmasks two distinct physical error sources previously concealed by local overfitting. First, manual selection introduces inherent target localization error, while coarse registration artificially satisfies these imprecise coordinates, global surface fitting exposes their true residual displacement. Second, the spatial discretization inherent to point cloud acquisition imposes finite sampling resolution and localized sensor noise, which subtly degrades fine geometric definitions surrounding the target landmarks. Consequently, the post-ICP metric increase does not indicate misregistration, but

rather confirms that the system successfully traded an over-constrained, noisy set of manual annotations for an authentic surface-wide fit.

5.2.2 Phase 2: Clinical Applicability and Inter-Operator Reproducibility

The quantitative registration metrics evaluated during the Phase 2 user study are summarized in Table 5.3, comparing the initial landmark-based coarse alignment against the final automated ICP refinement. Tracking performance is evaluated in terms of surface fitness and inlier RMSE, with overall results reported as mean ± standard deviation (SD).

Table 5.3: Quantitative alignment metrics of the user study evaluating the registration pipeline on the external knee surface ($n = 15$)

	Fitness (%)		**Inlier RMSE (mm)**	
Subject	**Coarse**	**ICP-refined**	**Coarse**	**ICP-refined**
1	100.0	100.0	3.491	0.899
2	100.0	100.0	4.241	0.896
3	100.0	100.0	3.851	0.897
4	100.0	100.0	3.972	0.896
5	100.0	100.0	3.158	0.898
6	100.0	100.0	2.812	0.896
7	100.0	100.0	3.181	0.892
8	100.0	100.0	3.019	0.909
9	100.0	100.0	3.276	0.908
10	100.0	100.0	2.984	0.899
11	100.0	100.0	3.793	0.891
12	100.0	100.0	3.153	0.889
13	100.0	100.0	2.733	0.873
14	100.0	100.0	3.416	0.887
15	100.0	100.0	3.289	0.887
Mean ± SD	100.0 ± 0.000	100.0 ± 0.000	3.358 ± 0.437	0.894 ± 0.009

The primary obstacle to reliable landmark-based registration is inter-operator subjectivity during soft-tissue target identification. Across 15 naive operators, manual coordinate selections varied by several millimeters, producing widely dispersed initial inlier RMSE values ranging from 2.733 mm to 4.241 mm (3.358 ± 0.437 mm). Automated point-to-plane ICP refinement neutralized this user-induced variability, compressing the error distribution into a deterministic envelope with an average inlier RMSE of 0.894 ± 0.009 mm.

This substantial drop in SD (from 0.437 mm to 0.009 mm) redefines the operational role of manual initialization. Rather than demanding expert-level placement from the clinician, the manual phase merely serves as a coarse spatial seed to place the registration parameters inside the non-linear basin of attraction of the ICP algorithm. Once trapped within this region, the continuous geometric topology of the knee surface dominates the trajectory of

the optimization, ensuring that final surgical accuracy remains largely independent of user proficiency. Importantly, this robustness was demonstrated across fifteen operators with no prior familiarity with the system, indicating that the workflow is resilient to variations in user execution. The consistent post-ICP accuracy despite differences in manual landmark placement demonstrates a high degree of operator independence, which is important for clinical translation.

This convergence behavior also explains why inlier RMSE, rather than the surface fitness score, serves as the decisive evaluation benchmark. Across all trials, fitness saturated at an unvarying $100.0 \pm 0.000\%$. Because the optical sensor captured a regional sub-patch completely circumscribed by the reference CAD volume under the configured correspondence limit, the fitness score encountered an absolute ceiling effect. While this confirms that the acquired data points remained within the bounding threshold, the metric possesses zero variance and cannot discriminate registration quality. Inlier RMSE instead provides the continuous distance metric necessary to quantify the physical convergence achieved by the automated optimizer.

5.2.3 Phase 3: Direct Surgical Target Validation on Femur Morphology

While soft-tissue external mapping provides excellent macro-localization, total knee arthroplasty demands precise registration directly on the target osseous structures. Consequently, the third validation phase transitions the framework onto bare bone morphology, specifically targeting the femur bone.

Table 5.4: Alignment metrics for femur models

Model	**Fitness [%]**	**Inlier RMSE [mm]**
Femur1	29.59	0.336
Femur1m	23.60	0.553
Femur2	31.39	0.398
Femur2m	11.81	0.555
Femur3	29.36	0.357
Femur3m	15.96	0.529

Across all six experimental evaluations, the registration pipeline consistently achieved sub-millimeter local targeting precision, with inlier RMSE values ranging from 0.336 mm to 0.398 mm for complete surface models and from 0.529 mm to 0.555 mm under partial exposure conditions. To correctly interpret the surface fitness metrics, it is necessary to distinguish between global mesh coverage and local surface alignment precision under strict correspondence constraints.

The surface fitness metric quantifies the percentage of points in the captured intraoperative point cloud that find a valid corresponding point on the pre-operative CT reference model within a maximum correspondence distance threshold of 0.5 mm. For the fully exposed,

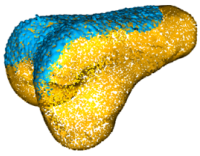
unmasked femurs, the fitness score clustered tightly between 29.36% and 31.39%. This ceiling is an inherent property of optical line-of-sight acquisition, whereas the pre-operative CT model is a complete 360° volumetric mesh, the intraoperative sensor captures only a single-view surface patch from the facing aspect, creating a partial-to-full registration scenario where roughly two-thirds of the reference model remains occluded. Under the masked conditions, physical occlusion by paper drapes restricted surface visibility even further, resulting in a marked drop in overall surface fitness to 11.81% to 23.60% . This configuration constituted a strictly partial-to-partial registration problem, where only a localized anatomical patch of the distal condyles was available to form surface correspondences.

The slight increase in inlier RMSE observed when transitioning from complete exposure (0.36 mm) to partial exposure (0.54 mm) is directly attributable to the loss of geometric leverage, removing the long lever arm of the femoral shaft reduces the spatial constraints available to lock down the six degrees of freedom during rigid transformation. Nevertheless, because the distal condylar region contains complex, non-planar curvature, the ICP algorithm successfully constrained the localized fit, maintaining inlier RMSE values well below 0.6 mm. This confirms that despite a severe reduction in total visible surface area, local alignment precision across the critical distal cutting zone remains high.

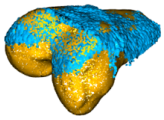
Comparing the masked femurs individually reveals a notable performance distinction between Femur 1m and the remaining masked models (Femur 2m and Femur 3m). In Femur 1m, the fitness score experienced a relatively modest drop from 29.59% to 23.60%. In contrast, Femur 2m and Femur 3m exhibited substantially sharper fitness reductions (31.39% to 11.81% and 29.36% to 15.96%, respectively).

This discrepancy can be explained by differences in coarse landmark initialization relative to the physical mask boundaries. Because initial coarse landmarks were selected on the unmasked femur reference point cloud prior to scanning, predicting the exact physical boundary of the paper drape was challenging. For Femur 2m and Femur 3m, one or more virtual coarse landmarks fell within regions that were physically occluded by the paper drape during scanning. This introduced landmark localization discrepancy during coarse point-picking, yielding a less optimal initial transformation that caused subsequent ICP optimization to converge toward a slightly degraded local minimum.

In contrast, the coarse landmarks selected for Femur 1m remained fully unobscured within the exposed anatomical zone. This enabled a superior coarse initialization, allowing ICP to converge to a tighter local alignment across the distal condyles, the primary region where partial-to-partial registration effectively takes place. This behavior is visually corroborated in Figures 5.1 to 5.3, where Femur 1m exhibits visibly tighter surface congruence along the condylar surface compared to Femur 2m and Femur 3m. Nevertheless, despite initialization variations, all masked specimens maintained inlier RMSE values near 0.55 mm, confirming that the pipeline retains robust targeting accuracy within the active surgical zone under constrained visibility.

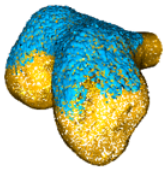


(a) Complete surface registration (Femur 1)

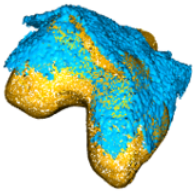


(b) Partial/masked surface registration (Femur 2m)

Figure 5.1: Qualitative visual comparison of the registration results between the fully exposed femur (a) and the clinically emulated partially masked femur (b) for model set 1

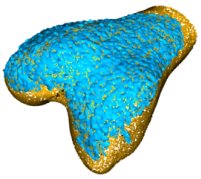


(a) Complete surface registration (Femur 2)

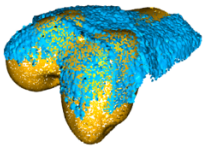


(b) Partial/masked surface registration (Femur 2m)

Figure 5.2: Qualitative visual comparison of the registration results between the fully exposed femur (a) and the clinically emulated partially masked femur (b) for model set 2



(a) Complete surface registration (Femur 3)



(b) Partial/masked surface registration (Femur 3m)

Figure 5.3: Qualitative visual comparison of the registration results between the fully exposed femur (a) and the clinically emulated partially masked femur (b) for model set 3

5.2.4 Current Limitations and Translational Challenges

Despite the strong performance metrics demonstrated across the validation phases, several technical limitations and translational barriers remain.

A primary workflow limitation is that the pipeline is not fully automated. The registration protocol relies on manual landmark selection to compute the initial coarse alignment. Although the subsequent ICP stage effectively mitigates human picking error and standardizes the final alignment outcome, this entry-level manual step requires direct operator intervention during intraoperative setup. This manual dependency increases setup time and introduces a potential source of failure if landmarks on the virtual CT model are difficult to identify or locate precisely on the physical target. Furthermore, all evaluations were conducted on static, 3D-printed phantoms using paper drapes to simulate surgical exposure.

While these setups provided a controlled and repeatable baseline, they do not capture complex intraoperative noise factors encountered in TKA, such as soft-tissue compliance and deformation, blood and fluid pooling within the surgical cavity, specular reflections from wet bone surfaces, and camera line-of-sight occlusions caused by surgical tools or the surgical team. Additionally, Phase 3 highlighted a structural sensitivity to partial surface exposure and draping boundary mismatches. Because the reference point cloud derived from the pre-operative CT scan represents the complete 360° bone geometry, physical draping intraoperatively restricts depth sensor access to a localized anatomical region. When coarse landmarks are assigned to regions that end up physically occluded by drapes, or when large portions of the reference mesh fall outside the visible field of view, the reduced geometric search space can cause the ICP optimization to converge toward a local surface minimum.

Finally, while the software architecture incorporates a continuous ICP module designed to adapt to minor object displacements, its tracking performance was not evaluated under active, dynamic limb manipulation. Conducting trials strictly under static conditions leaves the framework's real-time tracking latency and resilience to involuntary patient movement unverified.

5.3 Projector–Camera Calibration and Validation

The resulting intrinsic and distortion parameters obtained from the calibration are summarized in Table 5.5.

Table 5.5: Optimized projector intrinsic and distortion parameters.

Category	**Parameter**	**Evaluated Value**
Projector Intrinsics (K_{proj})	Focal length (f_x, f_y)	2634.19, 2631.77 px
	Principal point (c_x, c_y)	597.40, 656.82 px
Projector Distortion (D_{proj})	Radial (k_1, k_2, k_3)	0.0, 0.0, 0.0
	Tangential (p_1, p_2)	−0.0605, 0.0059
System Performance	Global Stereo RMSE	1.3803 px

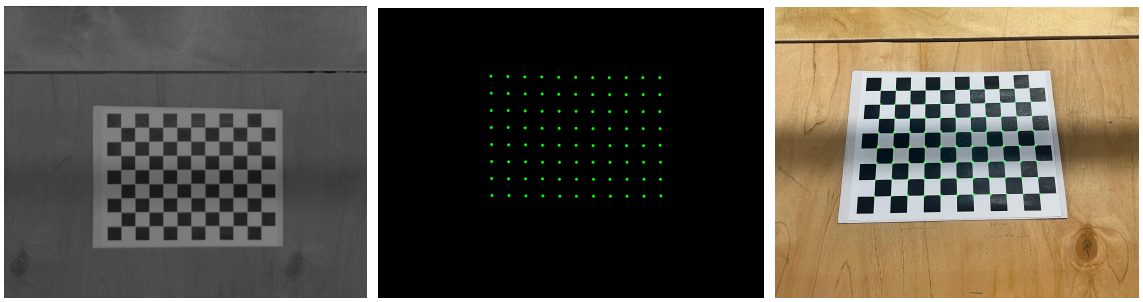
The global stereo RMSE converged to 1.3803 pixels. The rigid-body transformation mapping coordinates from the Ensenso camera reference frame to the projector frame is defined by the rotation matrix R and translation vector T (expressed in millimeters):

$$R = \begin{bmatrix} -0.99899283 & -0.00719239 & -0.04428992 \\ 0.00126086 & -0.99117937 & 0.13252118 \\ -0.04485240 & 0.13233186 & 0.99019015 \end{bmatrix} \tag{5.1}$$

$$T = \begin{bmatrix} 85.17404672 \\ -125.24576886 \\ 129.09629108 \end{bmatrix} \tag{5.2}$$

Computing the Euclidean norm of the translation vector yields a physical baseline distance of 199.02 mm between the optical centers of the Ensenso camera and the projector.

The stages of the closed-loop validation experiment are shown in Figure 5.4. First, the camera acquires an image of the physical calibration target to detect the reference checkerboard vertices (Figure 5.4a). Based on the calibrated extrinsic and intrinsic parameters, concentric circular targets are synthetically generated on the projector's display canvas corresponding to these detected 3D coordinates (Figure 5.4b). Finally, visual inspection of the resulting physical overlay (Figure 5.4c) demonstrated precise concentric alignment between the projected green circles and the physical vertices of the checkerboard squares across the entire projection field.



(a) Camera capture of target (b) Synthetic projection canvas (c) Physical projection overlay

Figure 5.4: Closed-loop back-projection validation

Fixing the Ensenso camera's factory intrinsics prevented the optimization routine from skewing the calibrated camera geometry to fit data artifacts, ensuring that residual optical and mechanical tolerances were mapped directly into the projector model and extrinsic transformation.

The convergence of the global stereo RMSE to 1.3803 pixels demonstrates high numerical

stability for an uncalibrated off-the-shelf projection unit. While radial distortion coefficients (k_1, k_2, k_3) were fixed to zero during optimization, the tangential distortion parameters ($p_1 = -0.06047783$, $p_2 = 0.00594048$) successfully absorbed minor physical decentering of the lens relative to the display chip. In addition, the near-identical focal lengths ($f_x \approx 2634.19$ px, $f_y \approx 2631.77$ px) confirmed a uniform pixel aspect ratio with negligible anamorphic distortion.

The evaluated baseline distance of 199.02 mm is consistent with the physical rig assembly. Furthermore, the close qualitative correspondence observed in the physical projection overlay confirms that the calibration pipeline successfully resolved the geometric transformation between camera and projector, establishing a reliable spatial mapping for real-time tracking and projection overlay.

5.4 SAR Navigation System

The quantitative performance of the end-to-end SAR projection and positioning pipeline across all three femoral specimens is summarized below. Table 5.6 presents the relative spatial distance and angular errors, while Table 5.7 details the absolute 3D residual displacement errors of the guide extremities following rigid registration.

Table 5.6: System-wide summary of end-to-end SAR projection and positioning errors across n=3 femoral specimens. Values are reported as mean ± SD

	Distance Error (Δd, mm)		**Angular Error ($\Delta\theta$, °)**	
Specimen	**SME**	**MAE**	**SME**	**MAE**
Femur 1	$+4.55 \pm 5.63$	5.50 ± 4.61	$+2.93 \pm 7.02$	5.16 ± 5.44
Femur 2	-0.33 ± 2.43	1.95 ± 1.37	$+1.27 \pm 6.02$	4.80 ± 3.59
Femur 3	$+4.20 \pm 2.81$	4.20 ± 2.81	$+3.59 \pm 3.93$	4.78 ± 2.15
Overall System	$+2.81 \pm 4.40$	3.88 ± 3.46	$+2.60 \pm 5.72$	4.91 ± 3.85

Table 5.7: Absolute guide point 3D residual displacement errors after rigid registration

Point	**Femur 1 (mm)**	**Femur 2 (mm)**	**Femur 3 (mm)**
G1	10.05	6.99	10.75
G2	32.55	16.59	13.70
G3	2.05	7.47	5.34
G4	3.85	18.75	11.98
Mean	12.12	12.45	10.44
Max	32.55	18.75	13.70

The experimental evaluation of the relative geometric metrics demonstrated an overall MAE of 3.88 ± 3.46 mm in spatial positioning and $4.91 \pm 3.85^\circ$ in angular orientation across the

three 3D-printed femoral specimens (Table 5.6). In the context of computer-assisted orthopedic surgery, conventional navigation standards target translational tolerances below 2.0 mm and angular alignment errors within $\pm 2.0^{\circ}$ to 3.0° across both coronal and sagittal planes to prevent premature implant wear and preserve joint kinematics [57]. While the aggregate system errors exceed these strict surgical thresholds, the primary conclusion from this evaluation is that the SAR pipeline successfully guides the surgical instrumentation to the correct anatomical region and target orientation. Given that this prototype represents a first-iteration benchtop implementation, the achieved gross spatial convergence confirms the feasibility of direct projection. Moreover, the relative performance recorded on Femur 2 (1.95 ± 1.37 mm MAE; -0.33 ± 2.43 mm SME; $4.80 \pm 3.59^{\circ}$ angular MAE) demonstrates that the underlying framework is capable of achieving sub 2 millimeter relative guidance when physical handling disturbances are minimized.

Complementing the relative metric analysis, the absolute point-based residual evaluation (Table 5.7) quantified the real-world spatial positioning discrepancies of the guide extremities after registering virtual planning coordinates to the benchtop femur reference frame. Across the three specimens, the mean 3D residual displacement was 12.12 mm for Femur 1, 12.45 mm for Femur 2, and 10.44 mm for Femur 3, with maximum deviations reaching 32.55 mm, 18.75 mm, and 13.70 mm, respectively.

While the relative distance deviations and absolute registered residuals evaluate fundamentally different geometric properties and are not directly comparable numerically, both metrics independently corroborate the same overarching conclusion. The system successfully directs surgical instrumentation toward the intended anatomical region and gross orientation, however, the observed deviations demonstrate that the current prototype remains far from the precision thresholds required for clinical adoption. Rather than indicating surgical readiness, both evaluations underscore that while the core concept of direct augmented projection is feasible, substantial refinement particularly in tracking stability, bone-guide mechanical fixation, and registration fidelity is essential before the pipeline can approach clinical navigation tolerances.

An analysis of the experimental workflow reveals that the larger deviations observed across both validations can be primarily explained by mechanical handling dynamics and benchtop setup constraints rather than fundamental computational failures of the tracking or projection algorithms. Both the physical surgical guide and the 3D-printed femur models featured rigid, low friction plastic contact interfaces, causing the guide to slide across the complex curved topology of the femoral condyles during manual alignment before the articulating support arm could be securely locked. Furthermore, the articulating mechanical arm exhibited notable joint stiffness and backlash, applying manual torque to lock its joints frequently induced an uncompensated rotational deflection that shifted the guide away from its intended alignment during final fixation. This mechanical drift was further compounded during the measurement phase as live tracking is paused to prevent user obstruction and the guide appa-

ratus from interfering with the registration, any subsequent physical contact such as aligning manual measurement tools against the guide extremities exerted forces on the support structure, resulting in contact induced micro-dislocations that were directly recorded as system error.

An inspection of the individual point measurements across the three specimens (Tables 5.8, 5.9, and 5.10) helps explain the distribution of the absolute residuals reported in Table 5.7. Rather than reflecting a uniform spatial offset across the entire assembly, the highest absolute residuals can be largely accounted for by localized discrepancies already evident at specific landmarks. In Femur 1, for example, the elevated residual at point G_2 (32.55 mm) is consistent with the localized shifts observed in its relative metrics, where that same point accounted for the largest individual distance deviation ($+14.1$ mm) and a $+19.26^\circ$ angular offset. Similarly, the higher residuals for points G_4 and G_2 in Femur 2 (18.75 mm and 16.59 mm) align with the notable angular shift recorded at G_4 ($+13.11^\circ$) and at G_2 ($+10.63^\circ$). These localized deviations indicate that particular guide extremities were subjected to greater physical displacement or manual reading disturbances during benchtop data collection, explaining why error spikes concentrate at those specific locations in both assessments.

The cumulative end-to-end positioning error represents the additive aggregation of uncertainties across every stage of the system pipeline. Upstream, the baseline stereo projector-camera reprojection error of 1.3803 pixels introduces an initial spatial ray-casting uncertainty at the operational working distance. Geometric discrepancies are further compounded by minor dimensional tolerances in CAD landmark extraction and structural wear on the physical guide interface, alongside residual surface-fitting errors from point-to-plane ICP alignment, which remains sensitive to local point cloud density, bone surface curvature, and sensor noise. Downstream, manually aligning the physical tool over the projected light contours relied on human visual servoing, where line-width dispersion and optical scattering across a guide surface exhibiting signs of wear from prior use inevitably limited placement precision. When combined with mechanical sliding of the guide prior to rigid fixation, clamping deflections, and inadvertent contact forces exerted during manual evaluation, these cumulative factors comprehensively account for the overall variance observed in the physical benchtop trials.

The experimental ground-truth validation protocol also possessed an intrinsic resolution limit that bounded measurement precision. Physical ink markings placed on the 3D-printed bones possessed a line stroke width of around 0.5 to 0.8 mm, introducing digitization uncertainty when associating visual ink centers with digital VTK cell-picked vertices (F_1, F_2, F_3). Furthermore, manual measurements taken with millimetric grid paper and vertical rulers had an inherent visual reading uncertainty. Consequently, in high-performing trials such as Femur 2, the measured deviation approached the accuracy limit of the physical measuring tools themselves. This limitation made it difficult to clearly separate the actual optical

projection and registration errors from the noise introduced by manual measurement.

Table 5.8: Surgical navigation verification: comprehensive planned vs. measured deviations for specimen set 1 (Femur 1).

Femur	**Guide**	d_{plan} **(m)**	d_{meas} **(m)**	Δd **(m)**	θ_{plan} (°)	θ_{meas} (°)	$\Delta\theta$ (°)
F1	G1	0.0674	0.0744	+0.0069	12.58	15.74	+3.16
	G2	0.0305	0.0406	+0.0102	40.93	60.19	+19.26
	G3	0.0654	0.0672	+0.0018	6.46	5.04	−1.42
	G4	0.0371	0.0348	−0.0023	51.48	47.83	−3.65
F2	G1	0.0394	0.0474	+0.0080	21.88	25.19	+3.30
	G2	0.0575	0.0656	+0.0082	20.33	32.51	+12.18
	G3	0.0391	0.0389	−0.0002	10.85	8.73	−2.13
	G4	0.0642	0.0628	−0.0014	26.88	24.24	−2.64
F3	G1	0.0530	0.0631	+0.0100	16.08	18.64	+2.56
	G2	0.0716	0.0857	+0.0141	16.19	24.31	+8.12
	G3	0.0557	0.0569	+0.0011	7.59	5.96	−1.63
	G4	0.0865	0.0847	−0.0018	19.61	17.71	−1.90

Table 5.9: Surgical Navigation Verification: Comprehensive planned vs. measured deviations for specimen set 2 (Femur 2).

Femur	**Guide**	d_{plan} **(m)**	d_{meas} **(m)**	Δd **(m)**	θ_{plan} (°)	θ_{meas} (°)	$\Delta\theta$ (°)
F1	G1	0.0781	0.0779	−0.0002	74.30	71.75	−2.56
	G2	0.1179	0.1208	+0.0029	31.16	34.04	+2.88
	G3	0.0774	0.0758	−0.0016	71.52	65.88	−5.64
	G4	0.1188	0.1192	+0.0004	30.38	31.62	+1.24
F2	G1	0.0807	0.0820	+0.0013	68.83	64.54	−4.29
	G2	0.1035	0.1028	−0.0007	36.13	41.15	+5.02
	G3	0.0829	0.0816	−0.0013	62.39	58.00	−4.40
	G4	0.1065	0.1018	−0.0047	34.34	37.87	+3.53
F3	G1	0.1010	0.1026	+0.0016	48.14	46.15	−1.99
	G2	0.0651	0.0687	+0.0035	69.48	80.11	+10.63
	G3	0.1019	0.0999	−0.0020	46.10	43.82	−2.28
	G4	0.0676	0.0645	−0.0032	62.73	75.84	+13.11

Table 5.10: Surgical navigation verification: comprehensive planned vs. measured deviations for specimen set 3 (Femur 3).

Femur	Guide	d_{plan} (m)	d_{meas} (m)	Δd (m)	θ_{plan} (°)	θ_{meas} (°)	$\Delta\theta$ (°)
F1	G1	0.0570	0.0663	+0.0093	76.13	72.32	−3.81
	G2	0.1054	0.1083	+0.0030	38.88	44.59	+5.72
	G3	0.0556	0.0593	+0.0037	85.30	81.98	−3.32
	G4	0.1045	0.1085	+0.0040	39.11	41.89	+2.79
F2	G1	0.0591	0.0663	+0.0071	69.33	72.45	+3.12
	G2	0.0932	0.0953	+0.0021	45.20	52.94	+7.74
	G3	0.0621	0.0638	+0.0017	63.06	67.00	+3.95
	G4	0.0952	0.0965	+0.0013	43.86	48.67	+4.80
F3	G1	0.0896	0.0952	+0.0057	38.14	41.56	+3.42
	G2	0.0682	0.0765	+0.0082	75.81	84.08	+8.28
	G3	0.0920	0.0925	+0.0006	37.01	39.37	+2.35
	G4	0.0706	0.0744	+0.0037	68.99	77.05	+8.06

Chapter 6

Conclusions and Future Work

The primary objective of this work was successfully achieved, demonstrating the feasibility of direct optical projection as an intuitive guidance medium for distal femoral resection in TKA. This prototype establishes a functional proof of concept, confirming that surgical cutting guides can be targeted onto the operative field without requiring invasive bone anchored tracking arrays or intramedullary canal drilling.

The experimental evaluation demonstrated that the multi-stage registration architecture provides a dependable foundation for patient-to-model alignment. On external knee surface the registration pipeline was successfully neutralized user subjectivity across 15 naive operators, condensing widely dispersed manual coarse alignments into a tightly bounded inlier RMSE of 0.894 ± 0.009 mm. When deployed directly on bare femur morphology, the pipeline consistently achieved sub-millimeter surface-fitting precision, yielding inlier RMSE values between 0.336 mm and 0.398 mm on complete bone surfaces, and maintaining values between 0.529 mm and 0.555 mm under simulated intraoperative draping occlusions. These outcomes confirm that surface-driven optimization can reliably register complex osseous morphology up to a satisfying degree even when optical line-of-sight access is restricted.

The physical benchtop validation of the complete SAR projection and positioning pipeline served as a functional proof of concept, successfully guiding the surgical cutting guide toward the correct anatomical region and gross orientation. Across the three evaluated specimens, the system achieved an overall relative positioning MAE of 3.88 ± 3.46 mm alongside an angular MAE of $4.91 \pm 3.85°$. In parallel, the absolute point-based evaluation across guide extremities yielded mean displacements between 10.44 mm and 12.45 mm, with maximum deviations reaching 32.55 mm. Although these absolute errors remain substantially larger than the standards mandated in clinical arthroplasty, the results demonstrate that the projection pipeline successfully conveys gross spatial trajectory and anatomical orientation, verifying the core potential of the direct optical guidance approach.

The registration framework, while effective on the segmented bone model, is not sufficiently robust against scene clutter, sterile drapes, and foreign objects located in the immediate vicinity of the surgical cavity; in future work, the registration pipeline must be improved to make it more resistant to such clutter. Live continuous ICP tracking was not tested in isolation under active dynamic conditions, meaning a dedicated validation methodology must be developed to rigorously evaluate this part of the system. Additionally, the physical validation methodology itself proved to be a major operational bottleneck. Measuring coordinates manually using millimetric paper grids and vertical rulers was exceptionally cumbersome, introduced an estimated human reading uncertainty, and made it difficult to reliably isolate optical projection inaccuracies from physical measurement errors. A new validation methodology must therefore be developed, as the current protocol introduces unnecessary errors due to measurement difficulty, susceptibility to unintended movement, and high reliance on human input that can lead to observer error. Finally, the physical guide support structure imposed severe mechanical constraints, which further hindered both validation and positioning. The articulating mechanical arm exhibited significant joint stiffness, mechanical backlash, and an awkward range of motion, making fine manual alignment over projected contours clumsy and causing unavoidable rotational shifts during final joint tightening and displacements during measurement. To overcome this limitation, an alternative holding arm or support fixture with higher mobility and stability must be developed in future iterations.

Appendix A

Results Tables

A.1 Raw Benchtop Coordinate Measurements

In this annex, the complete pre-operative ground truth coordinates and physical benchtop measurements (X, Y, Z) for all three 3D-printed femoral models ($n = 3$) are reported.

Table A.1: Femur 1: 3D coordinates of the femur (F) and guide points(G)

Dataset / Configuration	Point	X	Y	Z
Femur 1 – Ground Truth (Planned)				
Femur Points	F1	0.055	0.113	0.708
	F2	0.042	0.079	0.706
	F3	0.021	0.067	0.721
Guide Points	G1	0.043	0.056	0.674
	G2	0.035	0.136	0.708
	G3	0.052	0.056	0.676
	G4	0.051	0.137	0.680
Femur 1 – Measured				
Femur Points	F1	0.207	0.152	0.058
	F2	0.218	0.188	0.059
	F3	0.233	0.200	0.038
Guide Points	G1	0.227	0.210	0.100
	G2	0.227	0.128	0.084
	G3	0.215	0.210	0.091
	G4	0.217	0.130	0.083

Table A.2: Femur 2: 3D coordinates of the femur (F) and guide points(G)

Dataset / Configuration	Point	X	Y	Z
Femur 2 – Ground Truth (Planned)				
Femur Points	F1	-0.117	-0.033	0.014
	F2	-0.101	-0.036	-0.015
	F3	-0.037	-0.043	-0.018
Guide Points	G1	-0.104	-0.110	0.017
	G2	-0.025	-0.105	-0.002
	G3	-0.102	-0.108	0.026
	G4	-0.022	-0.104	0.007
Femur 2 – Measured				
Femur Points	F1	0.196	0.252	0.040
	F2	0.223	0.233	0.040
	F3	0.224	0.168	0.031
Guide Points	G1	0.181	0.219	0.109
	G2	0.217	0.147	0.096
	G3	0.174	0.216	0.103
	G4	0.212	0.145	0.090

Table A.3: Femur 3: 3D coordinates of the femur (F) and guide points(G)

Dataset / Configuration	Point	X	Y	Z
Femur 3 – Ground Truth (Planned)				
Femur Points	F1	0.034	0.030	0.100
	F2	0.051	0.031	0.073
	F3	0.104	0.024	0.074
Guide Points	G1	0.028	-0.023	0.080
	G2	0.106	-0.044	0.079
	G3	0.027	-0.024	0.089
	G4	0.106	-0.045	0.089
Femur 3 – Measured				
Femur Points	F1	0.188	0.199	0.029
	F2	0.212	0.180	0.036
	F3	0.212	0.129	0.026
Guide Points	G1	0.181	0.188	0.094
	G2	0.176	0.111	0.091
	G3	0.172	0.188	0.085
	G4	0.171	0.107	0.084

References

[1] Comprehensive Orthopaedics. Anatomy of the Knee. https://comportho.com/anatomy/anatomy-of-the-knee/, July 2016. [Online; accessed 2026-08-10].

[2] Knee Anatomy | OhioDance. https://ohiodance.org/dance-education/dance-wellness/knee-anatomy/, January 2013. [Online; accessed 2026-08-13].

[3] Meniscal Surgery. https://www.therapy-specialists.com/Injuries-Conditions/Knee/Surgery/Meniscal-Surgery/a~353/article.html. [Online; accessed 2026-09-05].

[4] U. F. O. Themes. The Knee. https://musculoskeletalkey.com/the-knee-10/, December 2016. [Online; accessed 2026-08-13].

[5] Zimmer, Inc. *Zimmer MIS Multi-Reference 4-in-1 Femoral Instrumentation Surgical Technique*. Zimmer, Inc., Warsaw, Indiana, USA, 2009. Rev. 7, Doc. No. 97-5967-002-00. Available: http://www.rpa.spot.pt/getdoc/03e37ea6-eb8d-409c-95e6-b227571b31dd/Nex_Gen.aspx [Accessed 2026-08-06].

[6] Robotic-Assisted Total Knee Replacement Surgery. https://uoanj.com/joint-replacement/robotic-assisted-total-knee-replacement-surgery/. [Online; accessed 2026-08-20].

[7] Total Knee Replacement Market Size & Share Analysis with Growth Trends | CAGR 4.33%. https://www.360researchreports.com/market-reports/total-knee-replacement-market-207384. [Online; accessed 2026-08-10].

[8] Jasvinder A. Singh, Shaohua Yu, Lang Chen, and John D. Cleveland. Rates of Total Joint Replacement in the United States: Future Projections to 2020–2040 Using the National Inpatient Sample. *The Journal of Rheumatology*, 46(9):1134–1140, September 2019. doi:10.3899/jrheum.170990.

[9] Austin F. Smith, Christian J. Eccles, Samrath J. Bhimani, Kevin M. Denehy, Rohat B. Bhimani, Langan S. Smith, and Arthur L. Malkani. Improved Patient Satisfaction following Robotic-Assisted Total Knee Arthroplasty. *The Journal of Knee Surgery*, 34(07):730–738, June 2021. doi:10.1055/s-0039-1700837.

[10] alessiadr. Robotics in TKA: state of the art and future perspectives. https://www.jointsjournal.org/2025/06/03/robotics-in-tka-state-of-the-art-and-future-perspectives/, June 2025. [Online; accessed 2026-06-08].

[11] Matthew A. Varacallo, T. David Luo, Ahmed Mabrouk, and Norman A. Johanson. Total Knee Arthroplasty Techniques. In *StatPearls*. StatPearls Publishing, Treasure

Island (FL), 2026. Available: http://www.ncbi.nlm.nih.gov/books/NBK499896/ [Accessed 2026-06-08].

[12] F. Canovas and L. Dagneaux. Quality of life after total knee arthroplasty. *Orthopaedics & Traumatology: Surgery & Research*, 104(1, Supplement):S41–S46, February 2018. doi:10.1016/j.otsr.2017.04.017.

[13] The Scientific Committee from the Personalized Arthroplasty Society (PAS) *. Standardizing definitions of the total knee alignment techniques: recommendations by the Personalized Arthroplasty Society. *EFORT Open Reviews*, 10(8):623–635, August 2025. doi:10.1530/EOR-2024-0120.

[14] Jonathan R. Danoff, Garrett Moss, Barthelemy Liabaud, and Jeffrey A. Geller. Total Knee Arthroplasty Considerations in Rheumatoid Arthritis. *Autoimmune Diseases*, 2013(1):185340, 2013. doi:10.1155/2013/185340.

[15] Jiwanjot Singh, Rachid Rassir, Inger N. Sierevelt, Oussama F.F. Van Geene, Maarten V. Rademakers, Anneke Spekenbrink-Spooren, and Peter A. Nolte. Total Knee Arthroplasty in the Post-Traumatic Knee: Revision Risks and Functional Outcomes Compared to Osteoarthritic Knees. A Report of Primary Procedures From the Dutch Arthroplasty Register. *The Journal of Arthroplasty*, 39(2):363–367, February 2024. doi:10.1016/j.arth.2023.08.050.

[16] David Johannes Haslhofer, Victoria Anelli-Monti, Peter Hausbrandt, Christian Kammerlander, and Antonio Klasan. No difference in surgical time and total theatre time between robotically assisted and computer assisted total knee arthroplasty. *Journal of Orthopaedics and Traumatology*, 25(1):52, November 2024. doi:10.1186/s10195-024-00798-8.

[17] Saverio Affatato, Paolo Erani, Maurizio Fersini, Vincenzo Contaldi, Anna Rita Terrizzi, and Antonio Licciulli. Preliminary In Vitro Wear Assessment of Ceramic Cemented Femoral Components Coupled with Polyethylene Menisci. *Materials*, 14(9):2112, April 2021. doi:10.3390/ma14092112.

[18] Spencer K Wilhelm, Jacob L Henrichsen, Matthew Siljander, Drew Moore, and Mark Karadsheh. Polyethylene in total knee arthroplasty: Where are we now? *Journal of Orthopaedic Surgery*, 26(3):2309499018808356, September 2018. doi:10.1177/2309499018808356.

[19] Engineering for sustainable development: executive summary - UNESCO Digital Library. https://unesdoc.unesco.org/ark:/48223/pf0000375634. [Online; accessed 2026-06-07].

[20] J. Gordon Betts, Kelly A. Young, James A. Wise, Eddie Johnson, Brandon Poe, Dean H. Kruse, Oksana Korol, Jody E. Johnson, Mark Womble, Peter DeSaix, J. Gordon Betts, Kelly A. Young, James A. Wise, Eddie Johnson, Brandon Poe, Dean H. Kruse, Oksana Korol, Jody E. Johnson, Mark Womble, and Peter DeSaix. 9.1 Classification of Joints - Anatomy and Physiology 2e | OpenStax. https://openstax.org/books/anatomy-and-physiology-2e/pages/9-1-classification-of-joints, April 2022. Book Title: Anatomy and Physiology 2e.

[21] Pallavi Juneja, Akul Munjal, and John B. Hubbard. Anatomy, Joints. In *StatPearls*. StatPearls Publishing, Treasure Island (FL), 2026. [Online; accessed 2026-06-09].

[22] Michael T. Hirschmann and Werner Müller. Complex function of the knee joint: the current understanding of the knee. *Knee Surgery, Sports Traumatology, Arthroscopy*, 23(10):2780–2788, October 2015. doi:10.1007/s00167-015-3619-3.

[23] Prathap Kumar J., Arun Kumar M., and Venkatesh D. Healthy Gait: Review of Anatomy and Physiology of Knee Joint. *International Journal of Current Research and Review*, 12(06):01–08, 2020. doi:10.31782/IJCRR.2020.12061.

[24] Jawad F. Abulhasan and Michael J. Grey. Anatomy and Physiology of Knee Stability. *Journal of Functional Morphology and Kinesiology*, 2(4):34, December 2017. doi:10.3390/jfmk2040034.

[25] Giles R. Scuderi and Alfred J. Tria. *The Knee: A Comprehensive Review*. World Scientific, 2010. Google-Books-ID: CdVpDQAAQBAJ.

[26] Anikar Chhabra, C. Curtis Elliott, and Mark D. Miller. Normal Anatomy and Biomechanics of the Knee. *Sports Medicine and Arthroscopy Review*, 9(3):166, September 2001.

[27] Andrew J. Price, Abtin Alvand, Anders Troelsen, Jeffrey N. Katz, Gary Hooper, Alastair Gray, Andrew Carr, and David Beard. Knee replacement. *The Lancet*, 392(10158):1672–1682, November 2018. doi:10.1016/S0140-6736(18)32344-4.

[28] Total knee replacement: Understanding patient-related factors. https://bcmj.org/articles/total-knee-replacement-understanding-patient-related-factors, November 2016. [Online; accessed 2026-06-10].

[29] David T. Felson. Osteoarthritis of the Knee. *New England Journal of Medicine*, 354(8):841–848, February 2006. doi:10.1056/NEJMcp051726.

[30] Leena Sharma. Osteoarthritis of the Knee. *New England Journal of Medicine*, 384(1):51–59, January 2021. doi:10.1056/NEJMcp1903768.

[31] Joern W.-P. Michael, Klaus U Schlüter-Brust, and Peer Eysel. The Epidemiology, Etiology, Diagnosis, and Treatment of Osteoarthritis of the Knee. *Deutsches Arzteblatt International*, 107(9):152–162, March 2010. doi:10.3238/arztebl.2010.0152.

[32] Walter Grassi, Rossella De Angelis, Gianni Lamanna, and Claudio Cervini. The clinical features of rheumatoid arthritis. *European Journal of Radiology*, 27:S18–S24, May 1998. doi:10.1016/S0720-048X(98)00038-2.

[33] Jin Kyu Lee and Choong-Hyeok Choi. Total Knee Arthroplasty in Rheumatoid Arthritis. *Knee Surgery & Related Research*, 24(1):1–6, March 2012. doi:10.5792/ksrr.2012.24.1.1.

[34] Hesham Saleh, Stephen Yu, Jonathan Vigdorchik, and Ran Schwarzkopf. Total knee arthroplasty for treatment of post-traumatic arthritis: Systematic review. *World Journal of Orthopedics*, 7(9):584–591, September 2016. doi:10.5312/wjo.v7.i9.584.

[35] Abdulaziz Alzarooni, Rana Muhammad Ali, Aqsa Aftab, Faizan Vaja, Muhammad Abdulvahab, Vijayalakshimi Rajmohankumar, Obed Amoako-Adjei, Sher Bahadur

Sunar, Gabriel G Hosu, Amena A Backosh, Bramaes Dahal, and Fadi A Jamaleddin Ahmad. A Comprehensive Review of Operative Considerations, Surgical Techniques, Outcomes, and Future Perspectives in Total Knee Arthroplasty. *Cureus*, October 2025. doi:10.7759/cureus.94345.

[36] Matthew A. Varacallo, T. David Luo, Ahmed Mabrouk, and Norman A. Johanson. Total Knee Arthroplasty Techniques. In *StatPearls*. StatPearls Publishing, Treasure Island (FL), 2026. [Online; accessed 2026-06-10].

[37] Muhammad Sohail, Salman Khalid, Muhammad Umar Elahi, and Heung Soo Kim. CAD-Based Analysis and Experimental Validation of Registration Errors in Imageless Total Knee Arthroplasty. *Journal of Clinical Medicine*, 15(6):2335, January 2026. doi:10.3390/jcm15062335.

[38] Xue Hu, He Liu, and Ferdinando M Rodriguez Y Baena. Automatic Markerless Tracking for Computer-Assisted Orthopaedic Surgery. pages 122–117. doi:10.29007/b6xt.

[39] Lei Chen, Changzhou Feng, Yunpeng Ma, Yikai Zhao, and Chaorong Wang. A review of rigid point cloud registration based on deep learning. *Frontiers in Neurorobotics*, 17, January 2024. doi:10.3389/fnbot.2023.1281332.

[40] Jiaolong Yang, Hongdong Li, and Yunde Jia. Go-ICP: Solving 3D Registration Efficiently and Globally Optimally. In *2013 IEEE International Conference on Computer Vision*, pages 1457–1464, Sydney, Australia, December 2013. IEEE. doi:10.1109/ICCV.2013.184.

[41] Paul J. Besl and Neil D. McKay. Method for registration of 3-D shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, page 586. SPIE, April 1992. doi:10.1117/12.57955.

[42] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, April 1992. doi:10.1016/0262-8856(92)90066-C.

[43] S. Rusinkiewicz. A symmetric objective function for icp. *ACM Transactions on Graphics (TOG)*, 38:1 – 7, 2019. doi:10.1145/3306346.3323037.

[44] Ki Hoon Kwon and Min Young Kim. Robust H-K Curvature Map Matching for Patient-to-CT Registration in Neurosurgical Navigation Systems. *Sensors*, 23(10):4903, January 2023. doi:10.3390/s23104903.

[45] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek. The Trimmed Iterative Closest Point algorithm. In *2002 International Conference on Pattern Recognition*, volume 3, pages 545–548 vol.3, August 2002. doi:10.1109/ICPR.2002.1047997.

[46] Felix Järemo Lawin, Martin Danelljan, Fahad Shahbaz Khan, Per-Erik Forssén, and Michael Felsberg. Density Adaptive Point Set Registration. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3829–3837, June 2018. doi:10.1109/CVPR.2018.00403.

[47] Georgios Dimitrios Evangelidis and Radu Horaud. Joint Alignment of Multiple Point Sets with Batch and Incremental Expectation-Maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1397–1410, June 2018. doi:10.1109/TPAMI.2017.2717829.

[48] Wei Gao and Russ Tedrake. FilterReg: Robust and Efficient Probabilistic Point-Set Registration Using Gaussian Filter and Twist Parameterization. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11087–11096, June 2019. doi:10.1109/CVPR.2019.01135.

[49] Andriy Myronenko and Xubo Song. Point Set Registration: Coherent Point Drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, December 2010. doi:10.1109/TPAMI.2010.46.

[50] Martin Magnusson, Achim Lilienthal, and Tom Duckett. Scan Registration for Autonomous Mining Vehicles Using 3D-NDT. *Journal of Field Robotics*, 24:803–827, October 2007. doi:10.1002/rob.20204.

[51] Wentao Yuan, Benjamin Eckart, Kihwan Kim, Varun Jampani, Dieter Fox, and Jan Kautz. DeepGMR: Learning Latent Gaussian Mixture Models for Registration. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 733–750, Cham, 2020. Springer International Publishing. doi:10.1007/978-3-030-58558-7_43.

[52] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast Point Feature Histograms (FPFH) for 3D registration. In *2009 IEEE International Conference on Robotics and Automation*, pages 3212–3217, May 2009. doi:10.1109/ROBOT.2009.5152473.

[53] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, June 1981. doi:10.1145/358669.358692.

[54] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast Global Registration. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 766–782, Cham, 2016. Springer International Publishing. doi:10.1007/978-3-319-46475-6_47.

[55] Dylan Campbell and Lars Petersson. GOGMA: Globally-Optimal Gaussian Mixture Alignment. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5685–5694, Las Vegas, NV, USA, June 2016. IEEE. doi:10.1109/CVPR.2016.613.

[56] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, Honolulu, HI, July 2017. IEEE. doi:10.1109/CVPR.2017.16.

[57] Taeho Jang and Kunwoo Lee. A novel registration method for total knee arthroplasty using a patient-specific registration guide. In *2012 12th International Conference on Control, Automation and Systems*, pages 1859–1862, October 2012. [Online; accessed 2026-06-12].

[58] Anna Gounot, Guillaume Dardenne, Baptiste Dehaine, Valérie Burdin, and Marion Decrouez. Patient-Specific Prediction of Total Knee Arthroplasty Surgical Exposure Using a Statistical Shape Model Augmented with Clinical Dataset*. In *2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 1–7, July 2025. doi:10.1109/EMBC58623.2025.11252778.

[59] Tiancheng Li, Peter Walker, Danial Hammoud, Liang Zhao, and Shoudong Huang. Partial-to-Full Registration based on Gradient-SDF for Computer-Assisted Orthopedic Surgery, October 2024. doi:10.48550/arXiv.2410.03078.

[60] Amy M. Morton, Bardiya Akhbari, Douglas C. Moore, and Joseph J. Crisco. Osteophyte Volume Calculation Using Dissimilarity-Excluding Procrustes Registration of Archived Bone Models from Healthy Volunteers. *Journal of orthopaedic research : official publication of the Orthopaedic Research Society*, 38(6):1307–1315, June 2020. doi:10.1002/jor.24569.

[61] Yuankun Wei, Ziyu Fu, Debing Zhang, Guangqi Wang, Chun Zhang, Xiang Xie, and Xiangyu Li. A Registration Method for Total Knee Arthroplasty Surgical Robot. In *2022 IEEE 6th Information Technology and Mechatronics Engineering Conference (ITOEC)*, volume 6, pages 535–538, March 2022. doi:10.1109/ITOEC53115.2022.9734563.

[62] Xue Hu, He Liu, and Ferdinando Rodriguez Y. Baena. Markerless Navigation System for Orthopaedic Knee Surgery: A Proof of Concept Study. *IEEE Access*, 9:64708–64718, 2021. doi:10.1109/ACCESS.2021.3075628.

[63] Xin Long, Zichao Zou, Shizhong Jiang, and Rongqian Yang. A Global Optimization Framework for Long Bone Registration in Computer-Assisted Orthopedic Surgery. In *2025 7th International Conference on Frontier Technologies of Information and Computer (ICFTIC)*, pages 1398–1404, December 2025. doi:10.1109/ICFTIC68075.2025.11324628.

[64] H. Bäthis, L. Perlick, M. Tingart, C. Lüring, D. Zurakowski, and J. Grifka. Alignment in total knee arthroplasty: A comparison of computer-assisted surgery with conventional techniques. *The Journal of Bone and Joint Surgery. British volume*, 86-B(5):682–687, 2004. doi:10.1109/cce.2016.7562647.

[65] Maedeh Kafian Safari, Morteza Mirzae, Behnaz Gheflati, Mostafa Sharifzade, Joel Zuhars, Sunil Rottoo, and Hassan Rivaz. Deep learning-based femoral reconstruction from intraoperative point clouds for enhanced knee arthroplasty registration. *International Journal of Computer Assisted Radiology and Surgery*, April 2026. doi:10.1007/s11548-026-03609-z.

[66] Z. Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334, November 2000. doi:10.1109/34.888718.

[67] Daniel Moreno and Gabriel Taubin. Simple, Accurate, and Robust Projector-Camera Calibration. In *Visualization & Transmission 2012 Second International Conference on 3D Imaging, Modeling, Processing*, pages 464–471, October 2012. doi:10.1109/3DIMPVT.2012.77.

[68] Bingyao Huang, Ying Tang, Samed Ozdemir, and Haibin Ling. A Fast and Flexible Projector-Camera Calibration System. *IEEE Transactions on Automation Science and Engineering*, 18(3):1049–1063, July 2021. doi:10.1109/TASE.2020.2994223.

[69] F. Sadlo, T. Weyrich, R. Peikert, and M. Gross. A practical structured light acquisition system for point-based geometry and texture. In *Proceedings Eurographics/IEEE VGTC Symposium Point-Based Graphics, 2005.*, pages 89–145, June 2005. doi:10.1109/PBG.2005.194069.

[70] Jamil Draréni, Sébastien Roy, and Peter Sturm. Methods for geometrical video projector calibration. *Machine Vision and Applications*, 23(1):79–89, January 2012. doi:10.1007/s00138-011-0322-3.

[71] Benjamin Close, Daniel B McCulley, and Bruce H Thomas. Late breaking results: ARPipes: Aligning Virtual models to their Physical Counterparts with Spatial Augmented Reality.

[72] Simon Willi and Anselm Grundhöfer. Robust Geometric Self-Calibration of Generic Multi-Projector Camera Systems. In *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 42–51, October 2017. doi:10.1109/ISMAR.2017.21.

[73] Ricardo R. Garcia and Avideh Zakhor. Geometric calibration for a multi-camera-projector system. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, pages 467–474, January 2013. doi:10.1109/WACV.2013.6475056.

[74] Daye Yoon, Jinyoung Kim, Jayang Jo, and Kibum Kim. Projection Mapping Method Using Projector-LiDAR (light detection and ranging) Calibration. In *2023 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 857–858, March 2023. doi:10.1109/VRW58643.2023.00271.

[75] Daisuke IWAI. Projection mapping technologies: A review of current trends and future directions. *Proceedings of the Japan Academy. Series B, Physical and Biological Sciences*, 100(3):234–251, March 2024. doi:10.2183/pjab.100.012.

[76] A. Grundhöfer and D. Iwai. Recent Advances in Projection Mapping Algorithms, Hardware and Applications. *Computer Graphics Forum*, 37(2):653–675, 2018. doi:10.1111/cgf.13387.

[77] Hirotaka Asayama, Daisuke Iwai, and Kosuke Sato. Fabricating Diminishable Visual Markers for Geometric Registration in Projection Mapping. *IEEE Transactions on Visualization and Computer Graphics*, 24(2):1091–1102, February 2018. doi:10.1109/TVCG.2017.2657634.

[78] Yuta Halvorson, Takumi Saito, and Naoki Hashimoto. Robust Tangible Projection Mapping with Multi-View Contour-Based Object Tracking. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 756–757, March 2022. doi:10.1109/VRW55335.2022.00229.

[79] Sora Hisaichi, Kiwamu Sumino, Kunihiro Ueda, Hidenori Kasebe, Tohru Yamashita, Takeshi Yuasa, Uwe Lippmann, Petra Aswendt, Roland Höfling, and Yoshihiro Watanabe. Depth-Aware Dynamic Projection Mapping using High-speed RGB and IR Projectors. In *SIGGRAPH Asia 2021 Emerging Technologies*, SA '21, pages 1–2, New York, NY, USA, December 2021. Association for Computing Machinery. doi:10.1145/3476122.3484843.

[80] Gaku Narita, Yoshihiro Watanabe, and Masatoshi Ishikawa. Dynamic projection mapping onto a deformable object with occlusion based on high-speed tracking of dot marker array. In *Proceedings of the 21st ACM Symposium on Virtual Reality Software and Technology*, VRST '15, pages 149–152, New York, NY, USA, November 2015. Association for Computing Machinery. doi:10.1145/2821592.2821618.

[81] Daiki Miyazaki and Naoki Hashimoto. Dynamic projection mapping onto non-rigid objects with dot markers. In *2018 International Workshop on Advanced Image Technology (IWAIT)*, pages 1–4, January 2018. doi:10.1109/IWAIT.2018.8369679.

[82] Parinya Punpongsanon, Daisuke Iwai, and Kosuke Sato. Projection-based visualization of tangential deformation of nonrigid surface by deformation estimation using infrared texture. *Virtual Reality*, 19(1):45–56, March 2015. doi:10.1007/s10055-014-0256-y.

[83] Kosuke Maeda and Hideki Koike. MirAIProjection: Real-time Projection onto High-speed Objects by Predicting Their 3D Position and Pose using DNNs. In *Proceedings of the 2020 International Conference on Advanced Visual Interfaces*, AVI '20, pages 1–5, New York, NY, USA, October 2020. Association for Computing Machinery. doi:10.1145/3399715.3399839.

[84] Vincent Gay-Bellile, Adrien Bartoli, and Patrick Sayd. Direct Estimation of Non-rigid Registrations with Image-Based Self-Occlusion Reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(1):87–104, January 2010. doi:10.1109/TPAMI.2008.265.

[85] Riverbank Computing. Pyqt6: Python bindings for the qt cross-platform application framework. https://riverbankcomputing.com, 2026. [Online; accessed 2026-07-08].

[86] Will Schroeder, Ken Martin, and Bill Lorensen. *The Visualization Toolkit*. Kitware, 4th edition, 2006.

[87] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018.

[88] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.

[89] Andriy Fedorov, Reinhard Beichel, Jayashree Kalpathy-Cramer, Julien Finet, Jean-Christophe Fillion-Robin, Sonia Pujol, Christian Bauer, Dominique Jennings, Fiona Fennessy, Milan Sonka, John Buatti, Stephen Aylward, James V. Miller, Steve Pieper, and Ron Kikinis. 3d slicer as an image computing platform for the quantitative imaging network. In *Magnetic Resonance Imaging*, volume 30, pages 1323–1341. 2012. doi:10.1016/j.mri.2012.05.001.

[90] Dassault Systèmes. *SOLIDWORKS 2024*. Dassault Systèmes, Vélizy-Villacoublay, France, 2024. Versão 2024.

[91] Calib.io ApS. Camera calibration pattern generator. https://calib.io, 2026. [accessed 2026-07-20].

[92] Jakob Wasserthal, Hanns-Christian Breit, Meyer T. Meyer, Maurice Pradella, Dan Hinck, Alexander W. Sauter, Tobias Heye, Daniel T. Boll, Joshy Cyriac, Yang Shan, and Martin Segeroth. Totalsegmentator: Robust segmentation of 104 anatomical structures in ct images. *Radiology: Artificial Intelligence*, 5(5), 2023. doi:10.1148/ryai.230024.

[93] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-ICP: A Globally Optimal Solution to 3D ICP Point-Set Registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(11):2241–2254, November 2016. doi:10.1109/TPAMI.2015.2513405.

[94] H. Luong, M. Vlaminck, W. Goeman, and W. Philips. Consistent icp for the registration of sparse and inhomogeneous point clouds. In *2016 IEEE Sixth International Conference on Communications and Electronics (ICCE)*, pages 262–267, 2016. doi:10.1109/cce.2016.7562647.