Group 07:

Co-Robotic Ultrasound Imaging of Breast Assisting Mammography

Students: Julian Brown, Kevin Wang, Yuxin(Ethan) Chen

Mentors: Dr. Emad Boctor, Dr. Web Stayman, Dr. Russell Taylor, Yixuan Wu

Abstract

This project aims to add Ultrasound imaging to the Mammography screening procedure in order to lower the percentage of patients that require further testing. To do this, a robotic system has been designed that can autonomously move across the patient's breasts after the X-Ray images are taken and acquire Ultrasound images of any regions of concern that may be present. The system allows the user to input the locations of one or more areas of concern, then moves an ultrasound probe to this location, and acquires ultrasound images. While more work is needed to develop this system into something that could be used on real patients, this project provides a solid base with which to continue to build

Introduction

This project is focused on building a robotic system to augment current mammography screenings with autonomously acquired ultrasound images. Mammography is a screening for breast cancer using X-Ray imaging. While it does greatly reduce breast cancer mortality, it is not without its problems[1]. Over 40 million screenings are performed in the US every year and around 6 million patients (15%) are called back for further testing[2]. This results in time and money lost on the part of the patients and doctors, not to mention the stress involved with the possibility of cancer.

Mammography and Ultrasound Imaging in combination have been shown to increase the accuracy of breast cancer screenings. Mammography is a highly specific test, but lacks sensitivity, while Ultrasound is highly sensitive, but lacks specificity. When used together they result in a test that is both highly specific and highly sensitive. A 2006 NIH study showed an increase in diagnostic accuracy from

[1] Mammography. Ganga Breast Care Centre. (n.d.). Retrieved February 27, 2022, from https://gangabreastcare.com/mammography.php
[2] *Breast cancer - statistics*. Cancer.Net. (2022, February 24). Retrieved February 27, 2022, from https://www.cancer.net/cancer-types/breast-cancer/statistics#:~:text=More%20women%20are%20diagnosed%20with,(in%20situ)%20breast%20cancer

0.78 for a traditional mammogram to 0.91 for a mammogram plus ultrasound examination[3].

While adding imaging of a different type to 40 million examinations a year could sound logistically daunting, a robotic system for augmenting mammography with ultrasound imaging would minimize the change to the current screening procedure. No dedicated ultrasound technicians would be needed, and the testing would be easily repeatable and accurate. It would save an immense amount of time and money for millions of patients and doctors every year.

Technical Approach

Calibration

There are four main parts for calibration, which are UR5 calibration, camera calibration, hand eye calibration and ultrasound probe calibration. A flowchart of the approach is shown in Figure 1.

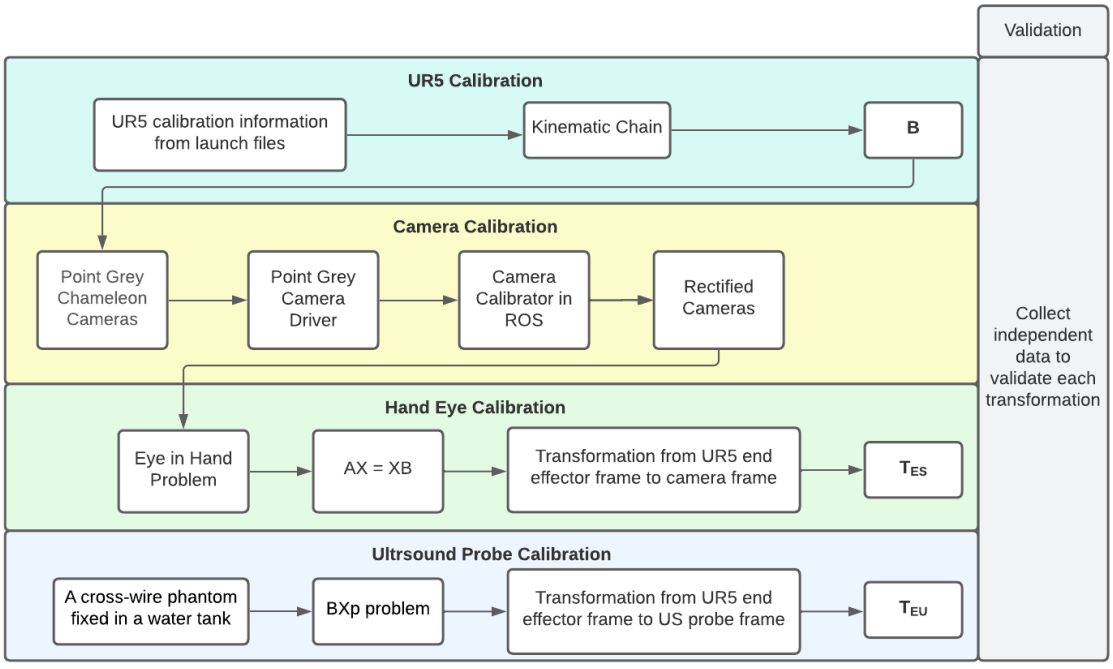

Figure 1: Flowchart Detailing Technical Approach Towards Calibration

[3] Wendie, A., Berg, M. D., & Jeffrey, D. (2008). Combined Screening with Ultrasound and mammography compared to mammography alone in women at elevated risk of breast cancer: results of the first-year screen in ACRIN 6666. JAMA, 299(18), 2151-2163..

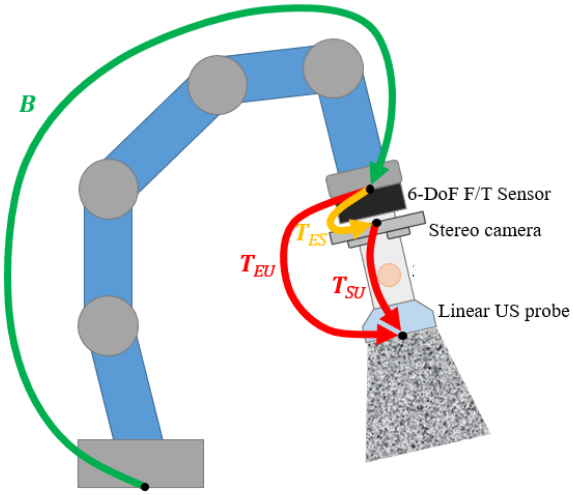

Figure 2: Transformations on setup

Camera Calibration

For camera calibration, a Point Grey Chameleon was calibrated to get rid of the distortion on the image edges. Point Grey camera driver was installed and used to connect the camera to ROS. The camera calibrator in ROS was implemented to automatically calibrate the cameras by using an 8*6 checkerboard with 25mm squares. Then, the rectification matrix and camera matrix were exported and applied through the image_proc ROS package to generate the rectified camera images.

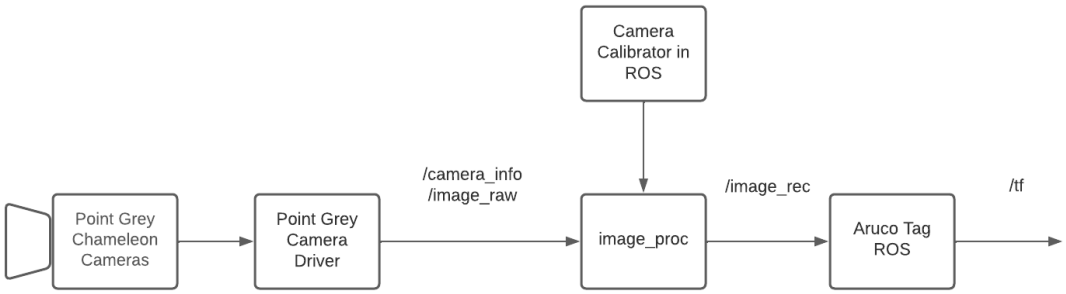

Figure 3: Flowchart of camera setup

Hand-Eye Calibration

For hand eye calibration, the goal was to find a 4 x 4 transformation matrix from the frame of the end-effector of UR5 to the camera frame. It was a hand-in-eye problem. An Aruco T ag was used as the marker. The Aruco Tag is made up of small 2D barcodes. The Aruco Tag ROS package was used to

recognize the location of the Aruco Tag and get the transformation from camera coordinate to the marker coordinate. The UR5 robot was moved to different poses where the camera was able to see the marker. The transformation from UR5 base frame to UR5 end effector frame and the transformation from camera frame to marker frame was recorded at each pose. Then, the transformation from UR5 end coordinate to camera coordinate was solved by establishing an AX=XB problem. Transformation data at twenty different poses was collected. Transformation data at another ten independent poses was also collected to validate this hand-eye calibration, which is shown in the results section.

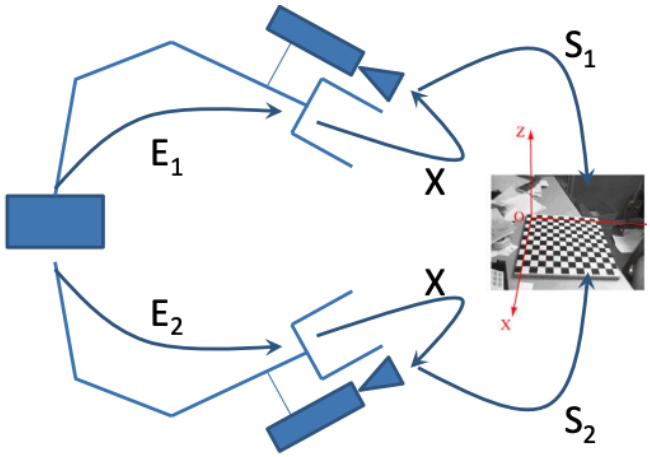

Figure 4: AX=XB hand eye calibration diagram[1]

AX=XB problem was established as shown in Figure 4. E represents the transformation from the base frame of UR5 to the end-effector frame of UR5. X represents the transformation from the end effector frame of UR5 to the camera frame. S represents the transformation from the camera frame to the 3D mark position. The equation of AX=XB is list as:

$$E_1 X S_1 = E_2 X S_2$$

$$X S_1 S_2^{-1} = E_1^{-1} E_2 X$$

$$AX = XB$$

To solve the AX=XB, the rotation part of X was solved first by using a polar decomposition method.

R_A is the rotation part of A; t_A is the translation part of A.

R_B is the rotation part of B; t_B is the translation part of B.

R_X is the rotation part of X; t_X is the translation part of X.

$$R_X = (M^T M)^{-1/2} M^T$$

$$where\ M = \sum_{i=1}^{N} \beta_i \alpha_i^T$$

$$\alpha_i = log(R_{Ai})$$

$$\beta_i = log(R_{Bi})$$

After the rotation part of X was found, it was then used to solve the translation part of X by using the least square method.

$$\begin{bmatrix} I - R_{A_1} \\ \vdots \\ I - R_{A_N} \end{bmatrix} t_X = \begin{bmatrix} t_{A_1} - R_X t_{B_1} \\ \vdots \\ t_{A_N} - R_X t_{B_N} \end{bmatrix}$$

Then the X is solved, which is the transformation from the UR5 end effector frame to the camera frame.

Ultrasound Calibration

For ultrasound probe calibration, the goal was to find a transformation matrix from the frame of the end-effector of UR5 to the ultrasound probe frame. A cross-wire phantom fixed in a water tank was used. The cross wire phantom is made of an intersection of two fishing lines as shown in Figure 5. The intersection can be seen as a point fiducial and can be observed under ultrasound. A cross wire phantom is very easy to set up and cheap. It has been proved in the previous research that it can acquire a very good calibration result. The transformation from the UR5 end-effector frame to the ultrasound probe frame was calculated by solving a BXp problem. Transformation data at 50 different poses was collected. Then transformation data at another 15 independent poses was also collected to validate this ultrasound calibration, which is shown in the result part.

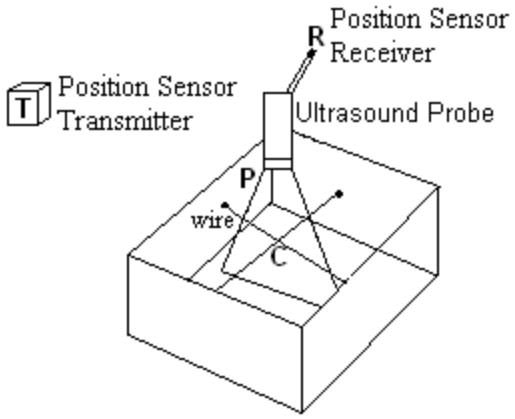

Figure 5: Cross wire phantom[2]

When collecting the data for calibration, the image depth of ultrasound was set to 5cm. The UR5 was moved to different positions to enable the point fiducial (the intersection of the fishing line) to be observed on the ultrasound image plane. Since the ultrasound image can only tell the 2D position of point fiducial on the ultrasound images. To get the transformation from end-effector coordinates of UR5 to the ultrasound image coordinates, the BXp problem was set up to solve this calibration problem, as shown in Figure 6.

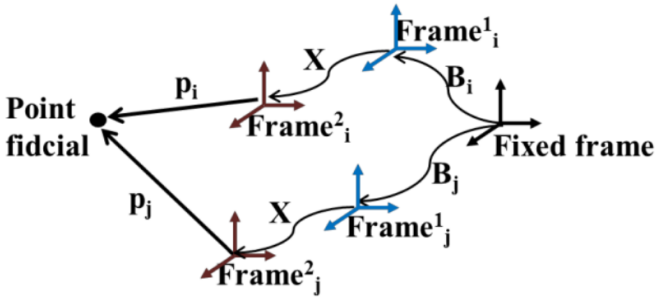

Figure 6: BXp calibration diagram [3]

B_iXp_i represents the point fiducial relative to the UR5 base frame, which is the fixed frame shown in the figure?, at the ith pose; B_i represents the transformation from UR5 base frame to the UR5 end effector frame; X is the transformation from the end effector of UR5 to the ultrasound probe frame.; p_i represent the location of point fiducial in the ultrasound image frame. Since the position of point fiducial relative to the UR5 base frame does not change. The BXp equation can be list as:

$$B_iXp_i = B_jXp_j$$

Using all the collected data to solve this BXp problem, the Gradient Descent algorithm was used [4]. Similar to the above equation, a cost function that need to minimize can be list as:

$$C_{BXp} = \left\| B_i Xp_i - B_j Xp_j \right\|_w$$

Then the Gradient Descent algorithm can be applied as [4]:

$$X_{s+1} = X_s \, exp(- \Delta t \alpha \widehat{\nabla C_{BXp}})$$

Where Δt is a small timestep; α is a learning rate; s+1 = sΔt.

An initial *X* was set as identity matrix. This gradient descent algorithm converged to the calculated *X* iteratively. Then the solved X is the transformation from end effector frame of UR5 to the ultrasound image coordinates.

Motion Planning

Robot motion planning with regards to ultrasound acquisition was split into two main motion segments: movement to a lesion site, and wobble about said lesion site. This approach is illustrated in Workflow 1.

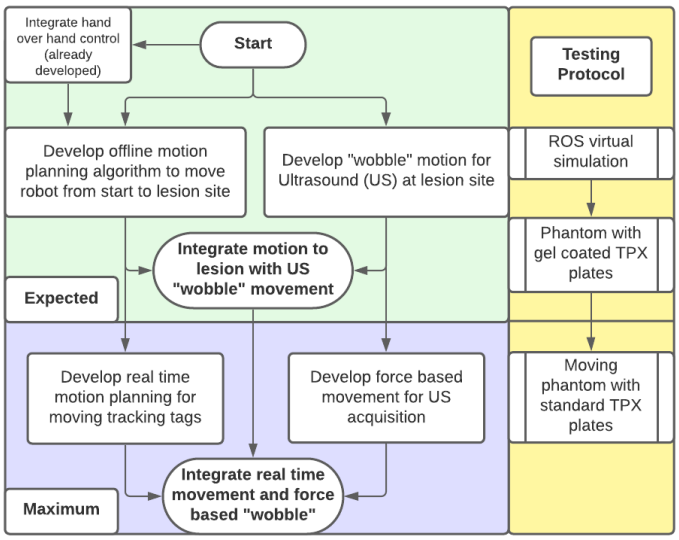

Workflow 1: Initial Plan for Co-Robotic Ultrasound Mammography Motion Planning Development

In the green box, is the plan to achieve the expected deliverables of motion. First, an offline robot motion planning algorithm for motion to the lesion site and the "wobble" motion were developed simultaneously.. These two prongs were integrated to complete the initial expected deliverable. In the

maximum deliverable outlined in the blue box, these motions were extended to account for more realistic scenarios; motion planning incorporates real time adjustments to account for patient movement, and the "wobble" motion incorporates force sensing to account for realistic US acquisition. The yellow boxes detail the testing protocol used. Expected deliverables were validated using a ROS virtual simulation then using a physical simulation consisting of a phantom in between acoustic gel coated plates. The maximum deliverable follows an identical testing protocol, with the addition of a moving phantom and standard TPX plates without the acoustic gel coating. The following sections will detail the specifics of each motion planning segment.

Dependencies

All motion planning and calibration were integrated on a linux system. In this system, various libraries are integrated into the ultrasound acquisition workflow. The overarching framework the co-robotic ultrasound mammography system uses is the Robot Operating System (ROS), an open source software development kit for robotics. This is a standardized platform for which robotic subsystems can be combined in a graph through processes known as nodes. The UR5 robot used for ultrasound acquisition is described virtually using the ur_modern_driver package. We used an Aruco Augmented Reality marker along with the Aruco ROS library to align our robot and target to a coordinate system. With the given target, motion planning was performed using RViz and MoveIt (C++). Finally, force control was integrated via a Robotiq FT-150 force sensor.

Ultrasound Probe

Robot description files were modified from the standard UR5 descriptions to account for the ultrasound probe attached to the UR5. Using the results from ultrasound calibration, an additional joint was added to the end effector link of the UR5. This allows for direct planning with respect to the ultrasound probe. Next, a conical obstacle was attached to the end effector to act as the ultrasound probe in planning to account for collisions between the probe and robot, as well as the probe and workspace.

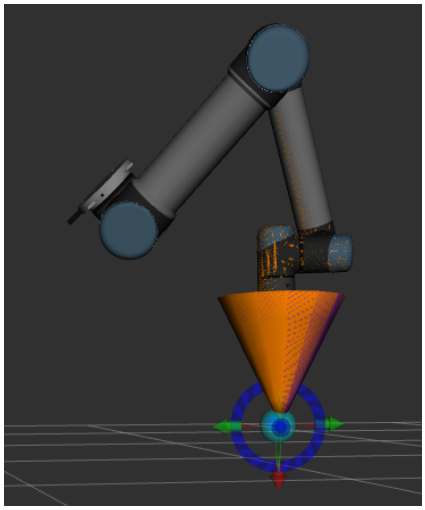

Figure 7: R-viz screenshot displaying the UR5 robot, including the conical obstacle attached to the end-effector to avoid collision with the ultrasound probe and camera.

Workspace

In addition to the modification of the UR5's robot description files, obstacles were added to the R-viz virtual workspace to prevent collisions with the table and the walls of the laboratory. Once these surroundings become more permanent, these obstacles can be edited and more can be added to represent the workspace even better.

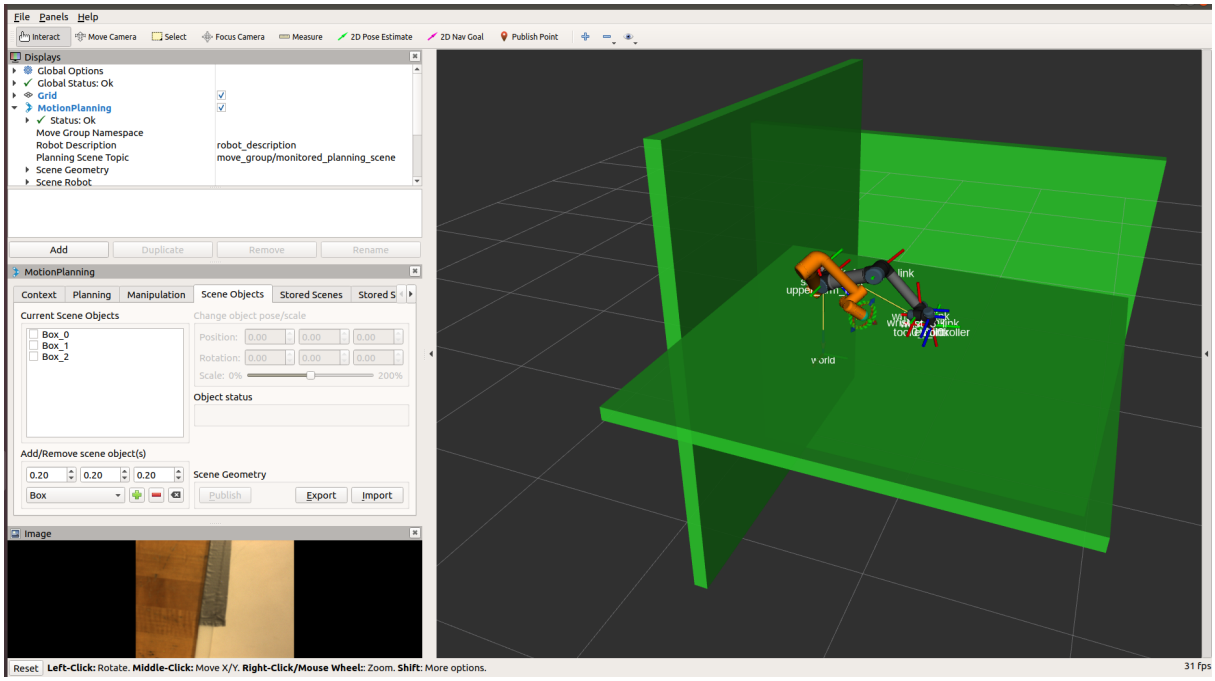

Figure 8: R-viz screenshot displaying the UR5 robot, including obstacles added to prevent collision with

the table and walls in the lab space.

Movement to a Point

With proper calibration, the Aruco Augmented Reality marker is able to be read with cartesian coordinates in relation to the base of the robot. This marker acts as the origin for the workspace, and a lesion location in relation to said origin is entered as an input to the motion planning algorithm. In order to move the UR5 to this entered position, the full pose of the probe is required - its position pointed vertically above the lesion site. The vertical orientation was extracted from the joint poses as the quaternion: [0.234713; 0.725081; -0.232557; 0.604222]. The probe position was determined as a translation from the detected Aruco marker and the inputted lesion location.

Force Sensing

Once the probe has reached a lesion site, it now must approach the plate until it comes into contact. In order to do this consistently, force data from a Robotiq FT-150 Force Sensor is used. The sensor is located between the end-effector of the UR5 and the ultrasound probe. Robotiq provides a ROS package alongside its sensors, allowing the data to be read into the computer and published onto a ROS topic. Our movement program subscribes to this topic. The algorithm is designed to lower the probe 1 mm at a time, checking the vertical force each time. Once that force surpasses a certain threshold, the probe stops, and the next segment of the motion control can begin.

Wobble

In order to acquire ultrasound images of a whole lesion, a wobble motion is performed at each lesion site. The wobble motion consists of a 30 degree motion to acquire ultrasound images from multiple angles. As such, a wobble about the lateral axis on the probe was programmed. To maintain a constant probe position, the wobble motion was divided into 3 degree increments. These orientations were calculated from the initial vertical orientation of the robot through the rotation matrix about the y axis shown below, where Θ = 0.0523598776 radians and - 0.0523598776 radians for left and right wobble respectively.

$$R_y(\theta) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix}$$

In the workflow, the robot acquires ultrasound data through a left followed by a right wobble.

Motion Planning Algorithms

The co-robotic ultrasound mammography system uses a Rapidly-Exploring Random Tree (RRT) method to plan its motion. RRT works by building a tree of nodes in the robot's configuration space. According to a certain probability called the goal bias, the algorithm will either attempt to add the goal position to the tree or add a randomly chosen point. Once the goal position is successfully added to the tree, the path is determined by tracing the nodes back from goal to initial. RRT is a simple and fast motion planning algorithm, and as the motion planning for this project is relatively simple with minimal obstacles, it was determined to be the best method for the job.

Results

Camera and Hand-Eye Calibration

The evaluation of camera calibration and hand-eye calibration was done together by moving the UR5 to 15 different poses where the camera can see the marker. The position of the marker relative to the base of the UR5 is calculated by using the following equation:

$$Position\ of\ marker\ relative\ to\ the\ UR5\ base\ frame\ =\ E\ \times\ X\ \times S$$

Where E is the transformation from the UR5 base frame to the UR5 end effector frame; X is the transformation from the UR5 end effector frame to the camera frame; S is the transformation from camera frame to the position of marker.

For the most ideal condition, since the position of the marker is fixed, the positions of the marker relative to the base of the UR5 should be exactly the same. However, since the calibration cannot be perfect, there was some error between the data collected at different poses.

The mean error of translation part is calculated by using the following formula:

$$mean\ error_{translation}\ =\ \frac{1}{N}\sum_{i=1}^{N}|t_i\ -\ \bar{t}|$$

Where t_i represents the translation part of the marker relative to UR5 base frame calculated at i^{th} pose. $\bar{t}$ is the average translation part of the marker relative to the UR5 base frame calculate for all the validation data.

The table below shows the mean error in each of the three axes.

Table 1: Mean error in each of the three axes, calculated using 15 different validation poses.

Mean error on X-axis of translation part	Mean Error on Y-axis of translation part	Mean error on Z-axis of translation part
0.0197 m	0.0113 m	0.0114 m

As seen in Table 1, the error was around 1-2 cm in each direction. The standard deviation of translation part is calculated by using the following formula:

$$std_{translation} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(t_i - \bar{t})^2}$$

Where t_i represents the translation part of the marker relative to UR5 base frame calculated at i^{th} pose. $\bar{t}$ is the average translation part of the marker relative to the UR5 base frame calculated for all the validation data.

The table below shows the standard deviation in each axis.

Table 2: Standard deviation in each of the three axes, calculated using 15 different validation poses.

Standard deviation on x axis of translation part	Standard deviation on y axis of translation part	Standard deviation on z axis of translation part
0.0038m	0.0026m	0.0016m

As seen in Table 2, the standard deviation was around 2-3 mm in each direction. The mean error of rotation part is calculated by using the following formula:

$$mean\ error_{roll} = \frac{1}{N}\sum_{i=1}^{N}|roll_i - \overline{roll}|$$

$$mean\ error_{pitch} = \frac{1}{N}\sum_{i=1}^{N}|pitch_i - \overline{pitch}|$$

$$mean\ error_{yaw} = \frac{1}{N}\sum_{i=1}^{N}|yaw_i - \overline{yaw}|$$

Where $roll_i,\ pitch_i,\ yaw_i$ represents the roll, pitch, yaw of rotation part of the marker relative to UR5 base frame calculated at i^{th} pose. $\overline{roll}$, $\overline{pitch},\ and\ \overline{yaw}$ are the average roll, pitch, yaw part of the marker relative to the UR5 base frame calculated for all the validation data.

The mean errors for roll, pitch, and yaw are shown in the table below.

Table 3: Mean errors for roll, pitch, and yaw, calculated using 15 different validation poses.

Mean error on roll of rotation part	Mean error on pitch of rotation part	Mean error on yaw of rotation part
0.0197 radian	0.0113 radian	0.0114 radian

As seen in Table 3, the mean error is around 0.01-0.02 radians for each euler angle. The standard of translation part is calculated by using the following formula:

$$std_{roll} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(roll_i - \overline{roll})^2}$$

$$std_{pitch} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(pitch_i - \overline{pitch})^2}$$

$$std_{yaw} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(yaw_i - \overline{yaw})^2}$$

Where $roll_i,\ pitch_i,\ yaw_i$ represents the roll, pitch, yaw of rotation part of the marker relative to UR5 base frame calculated at i^{th} pose. $\overline{roll}$, $\overline{pitch},\ and\ \overline{yaw}$ are the average roll, pitch, yaw part of the marker relative to the UR5 base frame calculated for all the validation data.

The table below shows the standard deviation for roll, pitch, and yaw.

Table 4: Standard deviation for roll, pitch, and yaw, calculated using 15 different validation poses.

Standard deviation on roll of rotation part	Standard deviation on pitch of rotation part	Standard deviation on yaw of rotation part
0.0246 radian	0.0138 radian	0.0143 radian

As seen in Table 4, the standard deviation was around 0.01-0.02 radians for each euler angle. All these mean errors and standard deviations are small enough to be acceptable, which proves that the camera calibration and hand eye calibration are accurate.

Ultrasound Calibration

To evaluate the ultrasound calibration, the robot was moved to 15 different poses where the ultrasound probe was able to see the point fiducial of cross wire phantom. The transformation B and point fiducial position in the ultrasound image coordinates was recorded. For evaluation, two metrics were defined. The accuracy metric can be defined as:

$$\sigma_{accuracy} = \left\| \sqrt{\frac{1}{N} \sum_{i=1}^{N} (B_i X p_i - C)^2} \right\|_2$$

The precision metric is calculated by using the following formula:

$$\sigma_{precision} = \left\| \sqrt{\frac{1}{N} \sum_{i=1}^{N} (B_i X p_i - \overline{BXp})^2} \right\|_2$$

Where B_i represents the transformation from UR5 base frame to the UR5 end effector frame; X is the transformation from the end effector of UR5 to the ultrasound probe frame.; p_i represent the location of point fiducial in the ultrasound image frame; c is the ground truth of the point fiducial position relative to UR5 base frame. The accuracy metric is the better evaluation metric, but in reality, the ground truth position can be achieved. Thus, the precision metric was used for the evaluation.

The result shows that the $\sigma_{precision}$ is 2.517 mm, which is acceptable. This proves that the ultrasound calibration is accurate.

Motion Planning

The final workflow for the co-robotic ultrasound mammography system is outlined in Workflow 2.

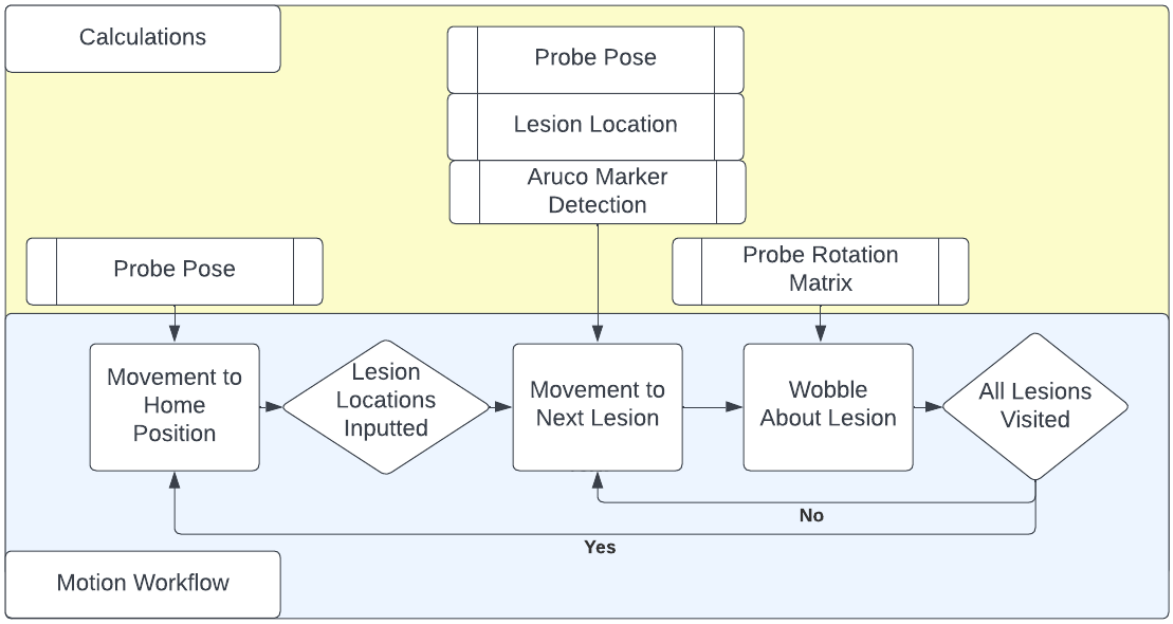

Workflow 2: Co-Robotic Ultrasound Mammography Ultrasound Acquisition Workflow

Initially, the UR5 arm moves to a home position that gives the robot a 30x25cm field of view, with the Aruco marker within this view. Next, lesion locations are inputted into the workflow. In this future, these locations would be detected via deep learning approaches. For each of the inputted lesions, the UR5 moves the US probe directly above the lesion location. Using force sensing, the probe is moved onto the lesion, until 2N of force is applied to the patient. Once in place a 15° and -15° wobble about the lateral axis of the probe are performed consecutively to acquire US data with a 30° range.

Wobble

Wobble for ultrasound acquisition was implemented as detailed in **Technical Approach.** The results of this implementation are illustrated in Figure 9.

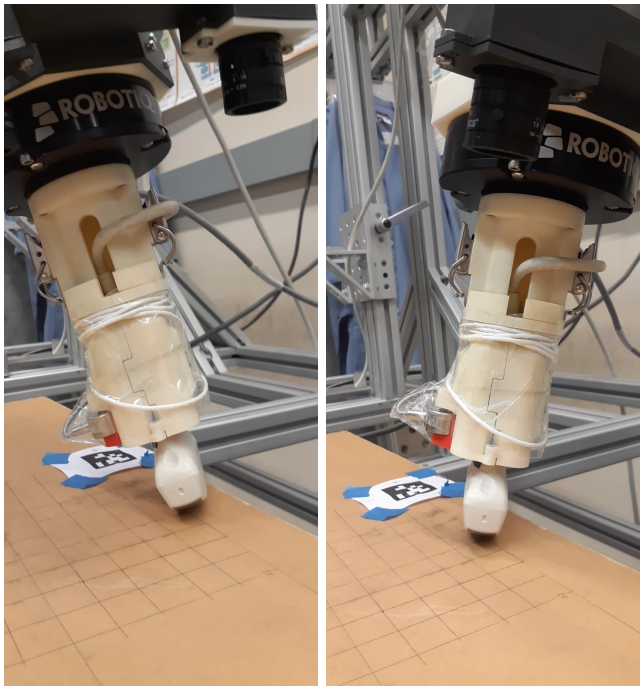

Figure 9: The ultrasound probe at each end of its 30 degree wobble motion.

Motion Planning Robustness

While the system is typically capable of successfully moving to the inputted points, approaching the plate through force sensing, performing a wobble, and returning to the initial position, there are times when a path is unable to be found or the robot undergoes a very complex path to its goal position. Some failures during the wobble motion are due to the robot being asked to extend out of its workspace. Another reason for wobble failure is that the robot sometimes must move upward and around, in order to get the full range of motion. Other times the reason for failure is unclear. More work needs to be done to improve the robustness of the system to ensure each movement can be successfully completed every time.

Acquisition of Ultrasound Images

Ultrasound images are acquired via the ultrasound probe and the Ultrasonix computer. Currently, that computer and the one operating the UR5 run independently, meaning that a technician would have to be on the Ultrasonix computer during system operation in order to take the pictures at the correct time. More work needs to be done to automate this aspect of the project, by getting the two computers to communicate to each other. This would allow the Ultrasound images to be taken automatically.

Conclusion

This project provides a foundation with which to continue the development of an autonomous system for augmenting Mammography screening with Ultrasound imaging. The system described over the course of this paper is able to accept the input of lesion site locations, move an ultrasound probe to that location, approach the plate with force-sensing, and perform a wobble motion to acquire ultrasound images at different angles. Future work will involve automating the acquisition of ultrasound images and to improvements towards the robustness of the UR5 motion planning.

Management Summary

Yuxin(Ethan):

Yuxin(Ethan) took charge of the camera calibration, hand-eye calibration, and ultrasound calibration, including the setup of the imaging validation and tests. Yuxin(Ethan) also contributed to the simple motion planning and real-time adaptation as the tracking tags moved.

During this semester, Ethan not only learned how to apply the camera calibration and hand eye calibration he learned in class to real project but also learnt how to do the ultrasound calibration. The team work experience on this project was also a great lesson for Ethan to understand how to collaborate with other engineers.

Julian:

Julian assisted Yuxin with calibration procedures and Kevin with the motion planning. He designed the user interface for the inputting of lesion site locations and augmented the motion planning with force sensing capabilities.

He learned how to control a UR5 in cartesian space, how to integrate sensor input into C++ code, how to calibrate an ultrasound probe to a robot, and more!

Kevin:

Kevin led the development of the robot workspace and motion planning algorithms, notably offline and online movement to a lesion, and ultrasound wobble. Kevin also assisted Yuxin and Julian in performing robot and ultrasound calibration.

Kevin learned how to interface with a physical robot using rViz and MoveIt to plan robot trajectories.

Proposed Timeline

Display original timeline along with actual

Table 5: Project Milestones and Dates

Milestone Name	Planned Date	Status
General Milestones		
Project Proposal Presentation	February 10	Complete
Project Proposal Document	March 1	Complete
Background Reading Presentation	March 10	Complete
Checkpoint Presentation	April 12	Complete
Demonstration Video	April 30	Complete
Calibration		
Camera Calibration	Feb 27	Complete
UR5 Calibration	Feb 27	Complete
Hand Eye Calibration	March 4	Complete
Ultrasound Probe Calibration	April 20	Complete
Robot Motion Planning		
Hand over Hand Control Integration	March 11	Removed
Offline Motion Planning to Lesion Site	April 15	Complete
“Wobble Motion” at Lesion Site	April 15	Complete
Maximum Deliverables		
Adaptive Motion Planning	April 29	Not Started
Force Sensitive “Wobble”	April 29	Complete
US - CT Registration	April 29	Not Started

Deliverables

Minimum:

- Calibration of the camera and robotic arm.
- Simple motion planning to land the ultrasound probe on a specific location on the compression plate. Offline (open loop) control, directing ultrasound probe to lesion from target.
- Setup of the imaging validation and tests.

Expected:

- A real-time interface to acquire camera and ultrasound images.
- With the developed real-time interface, demonstrate automatic ultrasound acquisition of a region of interest (whole lesion) with a known location with respect to the tracking tag on the compression plate.
- Documentation

Maximum:

- Demonstrate dynamic and real-time adaptation as the tracking tags move.
- Force control of the robot arm to ensure safety.
- Integration and registration with a second modality (Mammography, CT, or preoperative 3D ultrasound).

Of the nine deliverables listed above, seven of them have been completed over the course of this semester. All of the minimum and expected deliverables have been completed. Force control was the only maximum deliverable to be completed. Motion planning adaptation to a moving tracking tag has not been completed. The system is able to function successfully at different tracking tag locations in the workspace, but during a test, the camera no longer sees the tag, so the robot is unable to adapt to its movement. Integration with a second modality was also not completed. The next step of this project will be to continue to achieve the one last maximum deliverable: integration and registration with a second modality. This will involve acquiring mammography data alongside US data. Another direction for next steps would be to mount a camera that is near the station separated from the robot, as opposed to on the UR5 arm. This would allow the robot to maintain its full field of view through the entire operation, facilitating real time adaptation.

Acknowledgements

Thank you to each of our mentors, Dr. Boctor, Dr. Taylor, Dr. Stayman, and Yixuan, for their support and guidance!

References

List all references

[1] G. D. Hager and S. Leonard, “Hand-eye calibration - cs.jhu.edu,” Algorithms for Sensor-Based Robotics. [Online]. Available: https://www.cs.jhu.edu/~sleonard/lecture03.pdf. [Accessed: 01-May-2022].

[2] H. K. Zhang, A. Cheng, Y. Kim, Q. Ma, G. S. Chirikjian, and E. M. Boctor, “Phantom with multiple active points for ultrasound calibration,” Journal of Medical Imaging, vol. 5, no. 04, p. 1, 2018.

[3] F. Aalamifa, “Co-robotic Ultrasound Tomography: A New Paradigm for Quantitative Ultrasound Imaging”, Dept. of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, US, 2016, [Online]. Available: https://jscholarship.library.jhu.edu/handle/1774.2/40409

[4] M. K. Ackerman, A. Cheng, E. Boctor, and G. Chirikjian, “Online ultrasound sensor calibration using gradient descent on the Euclidean group,” *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

Technical appendices

For access to our code and instructions for installation and use of our ROS packages, please see our wiki page.