



RAM

● ROBOTICS
AND
MECHATRONICS

A NEEDLE GUIDANCE SYSTEM EMPLOYING ULTRASOUND REGION TRACKING FOR PRECISION BIOPSY PROCEDURES

B.N.E. (Benjamin) Meijerink

MSC ASSIGNMENT

Committee:

dr. F.J. Siepel
dr. V. Groenhuis, MSc
dr. ir. L.J. Spreeuwers

July, 2024

019RaM2024
Robotics and Mechatronics
EEMCS
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands

Abstract—An ultrasound-guided biopsy is a cost-effective and minimally invasive approach to taking biopsies from a subcutaneous lesion. These are primarily done by hand, leaving room for human error. Introducing a needle guidance system to aid the operator could enhance the accuracy and speed of the procedure, reducing time and costs. This paper introduces such an application, capable of tracking a lesion and manipulating the angle of the needle to assist during the biopsy. Experiments on a phantom and in vivo videos show that the tracking algorithm can track the target with deviations below 2 mm. The needle guide can accurately reach the lesion with a deviation of 0.6 ± 0.4 mm on ex vivo tests, compared to 1.1 mm to 2.9 mm in state-of-the-art research on in vivo tests. In conclusion, the proposed system can accurately target a lesion for biopsies while still being simple, portable, and cost-effective.

1 INTRODUCTION

1.1 Clinical background

Percutaneous core needle biopsy (CNB), guided by imaging techniques, is one of the best and most prevalent diagnostic methods for breast lesions and an alternative to biopsies requiring exposure to radiation or surgery [1]–[5]. A CNB performed on a breast can be seen in Figure 1. The needle is inserted to extract some tissue (a specimen). There are several benefits regarding percutaneous image-guided CNB compared to biopsies using surgery or stereotactic radiosurgery. Since it is less invasive, faster, and cheaper [6].




Fig. 1: Core needle biopsy on a breast [7].

Various medical imaging techniques can guide the needle by visualizing the needle and the target, like ultrasound (US), magnetic resonance imaging (MRI), and computer tomography (CT) [4], [8], [9]. Using US for imaging offers several benefits since the procedure typically can be done in real-time and is faster and cheaper than CT scans and MRIs. It also generates images without exposing the patient to ionizing radiation, which happens with CT scans. [4]. Nonetheless, CNB guided by ultrasonography is not perfect since multiple (3-4) specimens are often needed [10], [11]. This could be because the needle deviation during a breast biopsy is around 2.4 mm which could result in a specimen that does not contain part of the lesion [12]. This procedure

would become more appealing if the amount of necessary specimens could be reduced, obtained quicker, and obtained more accurately.




Fig. 2: The system consists of a US probe, needle, servo motor, and a 3D-printed needle holder, servo holder, and probe sleeve.

In the United States alone 1.6 million breast biopsies are performed each year [13]. As such, it would be beneficial if CNB with US guidance could be improved to reduce operating time and cost. Designing a system that could assist the operator in guiding the needle to the lesion could enhance the appeal of CNB.

1.1.1 State of the Art

The needle guidance tools that presently exist often consist of a physical tool that assists the operator in enhancing the accuracy, software that indicates the path of the needle, or a combination of the two [14]–[16]. Brattain et al. created an arm attached to the US probe that could assist the operator [14]. The arm keeps the needle in the probe’s scanning field and the projected needle path is indicated with sensors in the arm. There is no motor in the arm to assist with the angle control, only sensors to indicate the expected needle path. Tielens et al. show a needle guide that indicates the predicted trajectory of the needle [15]. Abayazid et al. introduce an application that tracks and controls a flexible needle with 3D US [16]. Research has also been done into contour tracking of a target on US images [17].

Additionally, research exists into investigating applications that combine these two and use a certain degree of automation to assist the operator, see Figure 3.

The following examples use ultrasound and robots for biopsies. However, these systems use big, complex robots or additional imaging forms, making the process time-consuming and the system not portable. Research into systems that solely rely on US, are simple and portable is less extensively documented. Maris et al. and Welleweerd et al. mention a robot arm with a 7-DOF serial manipulator with a needle guide using US [18], [19]. However, this application is not portable and still depends on MRI, making it expensive and time-consuming. Kronreif et al. discuss another US-guided biopsy robot, but this device also uses CT scans, is more expensive, and is big [20], [21]. Similarly, Kettenbach et al. propose the B-Rob I, a 7-DOF robot compatible with US- and CT-guided biopsies, and the



Fig. 3: Various state-of-the-art robots used for breast biopsies. A - Maris et al. & Welleweerd et al. [18], [19]; B - Kronreif et al. [20]; C & D - Kettenbach et al. [21], [22]; E - Jácomo-Zavaleta et al. [23]; F - Liang et al. [24].

B-Rob II a 4-DOF robot used for percutaneous interventions [21], [22]. Jácomo-Zavaleta et al. mention another biopsy robot using 3D ultrasonography [23]. Expanding on the 3D-US, Liang et al. mention a 3D-US-assisted autonomous biopsy robot. This application uses image thresholding to automatically find a lesion on the US scan [24]. Megali et al. mention a US-assisted needle guide robot using an 8-DOF robot arm where the operator indicates a target point and the needle's insertion point on the US frame [25]. Yang et al. discuss a US-assisted telerobot with a 6-DOF robot arm for remote-controlled biopsies [26].

1.2 Needle guidance for biopsies

This paper discusses an approach to automate part of the US-guided CNB procedure. The aim is to develop a simple and portable system that tracks a user-defined target on the US image. The target is the lesion that requires a biopsy. The operator can indicate the lesion with a target window. The projected needle path can be displayed in the frame to aid the operator.

Subsequently, the system uses a motor to manipulate the needle angle so the needle points to the lesion. The needle can then be inserted manually.

The motor can be controlled based on the target window's centroid (hereafter called the CoM) location. This technological development could assist the doctor in guiding the needle to the lesion in a more accurate and precise manner. This could reduce the possibility of wrong insertions, the patient's recovery time, the duration of the operation, and as a result, the cost of the operation.

Hence, the objective is to develop a US-guided CNB system that tracks the needle and the target in the body. The main goal is to determine how accurately a motorized needle guide attached to an ultrasound probe can be aligned with a lesion that is being tracked in real-time.

To achieve this goal, this research specifically investigates how a target lesion can be tracked on a live feed and how a motor can be controlled in a functional system such that the needle guide is steered toward the tumor. Research is also done to determine how a functional design for the application can be made. The tracking accuracy is evaluated for both in vivo and ex vivo situations. The biopsy accuracy is evaluated for ex vivo situations, compared to the required accuracy. Additionally, the advantages and disadvantages of the system are discussed.

First, the theoretical background will be discussed (Section 2), followed by the design method of the application and the experimental setup (Section 3). After this, the results of the experiments (Section 4) will be discussed followed by a discussion and conclusion (Sections 5 & 6).

2 THEORY

2.1 Ultrasound

Ultrasonography or ultrasound, is a medical imaging technique that can produce images of the inside of the body using high-frequency sound waves [17], [27], [28]. Despite creating a less accurate image, it is an attractive form of imaging since it is in real-time, inexpensive, and does not expose the patient to radiation, unlike CT scans and X-rays. Additionally, it is portable and does not require the patient to be put into a small space, as opposed to MRI scans.

Furthermore, using US during an operation like a biopsy is possible. This is not possible with a CT or PET scan and it is complicated with an MRI due to the magnetic field and the lack of space [29].




Fig. 4: B-mode US scan of benign lesions in a breast [30].

US images are created with acoustic waves with a frequency range typically between 2 and 18 MHz [27]. The waves are produced by supplying a current to piezoelectric crystals in the probe [31]. In turn, the returning waves cause the piezoelectric crystals to vibrate, creating a current that can be converted to an ultrasound image. For this image, the assumption is made that the speed of sound in the human body is uniformly 1540 m s^{-1} [32].

The transducer is placed on the skin so the high-frequency waves travel into the body. The echoes are used to distinguish between different tissues and boundaries, which generates images of the inside of the body in real-time. The tissues have differing properties like elasticity and density. When the waves reach the boundaries of tissues with different properties the waves are partially reflected to the transducer, the other part is scattered in other directions. The received waves are processed and, using the echos, the locations of the reflections are reconstructed. As a result, the boundaries become visible on the ultrasound image, making it possible to distinguish between objects. When different parts have similar acoustic properties it can be

hard/impossible to see them on a US scan. In the case of a lesion that is not detectable on US video but can be detected on MRI or CT, it is called a US/sonographically occult lesion [33], [34]. Additionally, some parts of the body reflect a lot of the acoustic waves to the transducer, other parts conduct the waves and return barely any waves. As a result objects on ultrasound can be classified into three categories: anechoic, hypoechoic, and hyperechoic [35].

Anechoic objects do not scatter the sound waves, because they are homogenous and conduct the waves well. Since no echos are reflected for these parts they appear black on the ultrasound image. Think of objects filled with fluids like cysts and veins. Hypoechoic parts reflect some echoes and as a result, they appear darker than surrounding tissues on the ultrasound image. This is the case for tumors and lymph nodes. The hyperechoic category consists of parts that reflect most of the waves, resulting in relatively bright parts on the ultrasound image. Hyperechoic objects are bones, fat, and lungs since the lungs contain air.

The most prevalent way to use US in medical imaging is the B-mode or brightness mode, seen in figure 4 [36]. This mode shows a cross-sectional image of the different tissues and their boundaries. Each boundary and tissue type has different properties and will return a different echo of the US waves. The echos are shown on the US image at the point corresponding with the point in the body where the echo originates. A grayscale value represents the strength of the echo. The brighter the point the stronger the echo. This results in a black-and-white image that shows a cross-section of the body in the probe's plane.

2.2 Software

This section discusses the theory behind the algorithms used in the application, (Section 3).

The software for this application consists of three parts. First, the US images from the transducer are retrieved. This will be represented in a live video feed. Second, the images are processed to track a user-defined window. Third, a motor steers the needle to the CoM of the window. In the optimal case, the CoM of this window aligns with the CoM of the lesion. At this point, the operator can insert the needle when a stable tracker position has been reached.

The software created for the tracking consists of Gaussian blur to compensate for noise and the tracking is done by the Mean Shift and CamShift algorithm. The performance of these two will be compared to see which works better. The angle is controlled through a transformation matrix which converts the target's pixel coordinates in the frame of the US transducer to coordinates in the frame of the needle on the motor. With these coordinates the required angle can be calculated using trigonometric identities.

2.2.1 Mean Shift

Mean Shift is a clustering technique to find the CoM of a density function [37]. It is widely used in image processing for segmentation and tracking.

The working principle of Mean Shift, shown in Figure

5, involves a window or kernel that is iteratively shifted towards the maximum density or the CoM of the data distribution of the confidence map within the image. In the case of an image, each pixel is seen as a data point, located in a feature space. The distribution of these points is described by the density function.

In the case of tracking a target on a US video, a window (hereafter referred to as the region of interest (ROI)) is selected. The ROI is a part of the US image. The intensities or grayscale values in the ROI can be represented by a histogram which serves as a probability density function (PDF).

Based on the histogram, a mask or confidence map is created, indicating the probability that each pixel in the image is part of the ROI.

The density of high confidence points is calculated and the difference between the CoM of the ROI and the CoM of the confidence points is determined with the Mean Shift vector. This vector indicates the direction where the increase in density is the greatest. The ROI is then adjusted based on the Mean Shift vector. Once the CoM of the window corresponds with the highest density of points, the target is found.

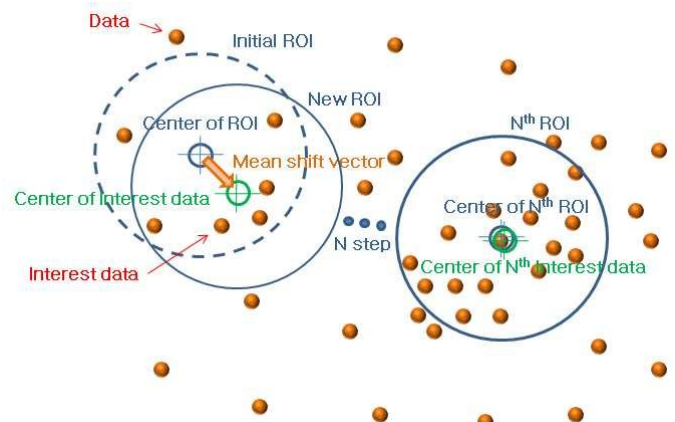


Fig. 5: Schematic of the Mean Shift algorithm [38]. It looks for the CoM of the pixels within the ROI.

2.2.2 CamShift

One limitation of Mean Shift is that the ROI is not variable in size. So if a target changes in size, due to a different positioning of the US transducer, the ROI cannot change with it. As a result, it might not indicate the CoM of the target but another place on the target. So, an enhancement of the Mean Shift algorithm is proposed, namely CamShift (Continuously Adaptive Mean Shift) [38]–[41]. It works similarly to the Mean Shift algorithm, but the ROI is adaptable in size. So if a target increases or decreases in size, the ROI size will adapt. This could increase the chance of the ROI's CoM corresponding with the target's CoM. The downside is that it is slower and can react more to outliers. A schematic is shown in Figure 6. In this case, the ROI is shaped so that the data fits better within the ROI.

2.2.3 Gaussian blur

Gaussian blur is a technique within computer vision used for image filtering and processing [42]. A Gaussian kernel is




Fig. 6: Schematic of the CamShift algorithm [38]. It looks for the CoM of the pixels within the ROI. The ROI is adaptable in size to better fit the data.

convolved with an image, which results in a smoother and blurred version of the original image. This could be useful when an image contains noise that negatively impacts the used computer vision algorithm. This can make a process, like tracking, more robust.

2.2.4 Angle control

The angle control is used to guide the needle to the target using a servo motor. Essentially two calculations are necessary for this. The coordinates need to be defined in the coordinate frame that has the needle guide as its origin. After this, a trigonometric identity has to be applied to calculate the angle.

To define the coordinates of the target in the frame of the servo the following formula can be used:

$$P^s = H_t^s P^t \quad (1)$$

Where P^s is the coordinate vector represented in the servo coordinate frame: $P^s = \begin{bmatrix} x_p^s \\ y_p^s \\ 1 \end{bmatrix}$ and P^t is the coordinate vector represented in the coordinate frame of the transducer:

$P^t = \begin{bmatrix} x_p^t \\ y_p^t \\ 1 \end{bmatrix}$. H_t^s is the 3×3 transformation matrix. The standard form is:

$$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$$

The 3×3 matrix can facilitate scaling, shearing, and rotational and translational transformations in the 2D space. In this matrix c and f are translational transformations, signifying the displacement of the origin. In this case, the displacement of the origin of the servo coordinate frame with regards to the origin of the transducer coordinate frame. a , b , d , and e factor in the scaling and shearing. They can also indicate rotations if these are involved. The bottom row is always the same.

3 METHOD

3.1 Design

The application has three main components: The biopsy needle, the US transducer, and the RC servo motor. A design has been created that combines these components.




Fig. 7: The CAD design of the needle guidance system. It consists of the top part of the sleeve for the probe (red), the bottom part of the sleeve connected to the servo holder (yellow), and a needle guide on top of the servo motor (green). The servo, probe, and needle are visible in the CAD design for reference.

The transducer and servo (through an Arduino) are connected to a laptop. The design ensures that the needle has two degrees of freedom. One translational degree of freedom, for insertion, and one rotational degree of freedom in the US transmission plane so the needle is visible on the US image. Movement in the other two rotational and translational planes is restricted. The hypothesis is that this movement restriction paired with the angle control could reduce the deviation from the needle to the lesion's CoM.

3.1.1 Mechanical design

Three parts are needed for the device. These are designed in SolidWorks. The proposed design can be seen in Figure 7, and the realized design can be seen in Figure 2. The parts are 3D-printed using tough PLA. The first is a part that guides the needle (green). This part is placed on top of the servo motor.

The needle holder can be attached to the top of the servo with a screw. It has a hollow conical part that fits exactly around the rotating part of the servo to prevent slipping. A hole is also needed in the front. This hole is placed off-center, away from the probe. As a result, the operating area of the application in the scanning field is bigger compared to a needle holder with a hole in the center.

A 14 gauge needle is used, which corresponds to a 2.1 mm outer diameter. Since the 3D printer makes the holes slightly smaller than designed a 2.1 mm drill bit is used. Which guarantees that the holes have the same diameter as the

needle. This ensures the needle can move forward and backward while restricting movement in the other directions. Having two partitions further ensures the needle can only move along one axis.

The other two parts are a sleeve encompassing the transducer (red and yellow) and a box for the servo (yellow). The box is made of four walls with one wall having a cutout for the wires. A ridge with holes at the back ensures the servo is fixed in place by screwing the box to the motor. The sleeve is made to fit snugly around the transducer to ensure it is not too thick. The transducer sleeve is printed in two parts that can be fixed together with screws.

The servo box and sleeve are combined so the needle is in the scanning field of the transducer and can be seen on the US scan. It should not enter the field from above with a downward angle or below with an upward angle. In such a case the needle would only be partly visible in the scanning field. As a result, the servo is located lower than the probe.

Additionally, it needs to be strong enough to handle the forces applied to it. So it should not be too brittle or thin. There also needs to be enough space between the servo and the transducer so the range of motion of the servo and needle is not limited.

Despite the size and location of the servo (holder) and the holes for the screws, the handling is still comfortable. A limitation of this device is that the needle cannot access all parts within the scanning field, but the device can be moved to ensure the target comes within reach of the needle.

A second motor could be added to increase the operating range of the application within the scanning field. However, this increases the complexity of the mechanical design and algorithm. It also makes the design more bulky. As a result, the probe would be harder to maneuver and it could be more challenging to examine certain spots with a bigger and bulkier application.

3.2 Software

The software consists of three parts: the US software retrieves the US frames from the US scanner and shows the US live feed, the tracking algorithm, and the needle-guiding algorithm. A flowchart, showing an overview of the code, can be found in Figure 8.

The first part is part of the US scanner and can be used as is. The tracking software is used to select and track a target and the needle-guiding algorithm points a needle to the CoM of the target using a servo motor.

The tracking algorithm consists of two parts: selecting and tracking the target. The user indicates the target using a mouse. This will be done using a rectangular region of interest (ROI) with a cross indicating the CoM. It is rectangular and not a contour of the lesion because the tracking algorithm depends on the intensity of the target and no geometric information is taken into account. So if the lesion deforms the ROI would not be accurate anymore. If the geometric information is accounted for, an ROI window could be made to fit the lesion's shape. Extra measures should then be taken to ensure that the ROI can adapt to deformations of the lesion.

3.2.1 Target tracking

The tracking part should be (close to) real-time to make the algorithm run smoothly. The bigger the delay is, the slower the operator should move the transducer and the longer they need to wait for the needle to be guided to the lesion. A delay of around 100 ms is acceptable, according to Ries et al. [43]. The paper mentions a latency smaller than 114 ms was allowed. So algorithms that are computationally very expensive can not be used.

One thing to consider is that no moving objects are in the frame. This is because there is barely any relative motion when moving the transducer, apart from differences in the deformation of the different tissues. Because of this, conventional algorithms that rely on background subtraction or relative movement are not useful for this application [44]. Additionally, the different tissues, like bones, tendons, tumors, and veins, deform differently. Some other challenges also arise regarding the tracking of the target.

Firstly, the US feed is noisy, so the algorithm needs to be robust to noise.

Secondly, the human body is soft, meaning that the target can be deformed when the transducer is moved. The target can also disappear from the frame when the transducer is moved. So the algorithm should be robust to deformations and occlusions.

The tracking of the target is done using Mean Shift and CamShift [37]. The explanation for these algorithms and the Gaussian blurring applied to the frame can be found in Section 2.

One limitation of this algorithm is that the skin in the top part of the frame is hypoechoic. So, if a lesion has a similar hypoechoic appearance it can pose problems when it is close to the skin or disappears from the frame. In this case, the tracker could mistake the skin for the target. By setting the top part of the confidence mask to zero confidence this problem can be resolved since it will not be seen as a target. Additionally, parts in the frame can have similar intensity values as the targeted lesion, so they will also be marked as high confidence. Sometimes these parts are other objects, but sometimes they can also be specks of noise. These smaller specks can be removed or mitigated using Gaussian blurring to increase the robustness of the tracker.

With this mask, the Mean Shift algorithm can find the confidence peak, which should correspond to the target's CoM. The initial window for the Mean Shift function is the user-defined ROI location. So even if other points in the frame have high confidence the function will track the correct target. This way the ROI can be tracked in the live feed. The downside of this method is that the 2D image is reduced to a 1D grayscale representation. So information about geometry and edges is lost. As long as the lesion is relatively homogeneous in color and has some contrast with the surrounding tissue this should not be a problem.

Since the Mean Shift algorithm has a fixed window size this could pose problems when the lesion size changes on the US image. Because of this, CamShift will also be tested to compare the performance. As explained in Section 2, this algorithm has a variable window size, which could help get

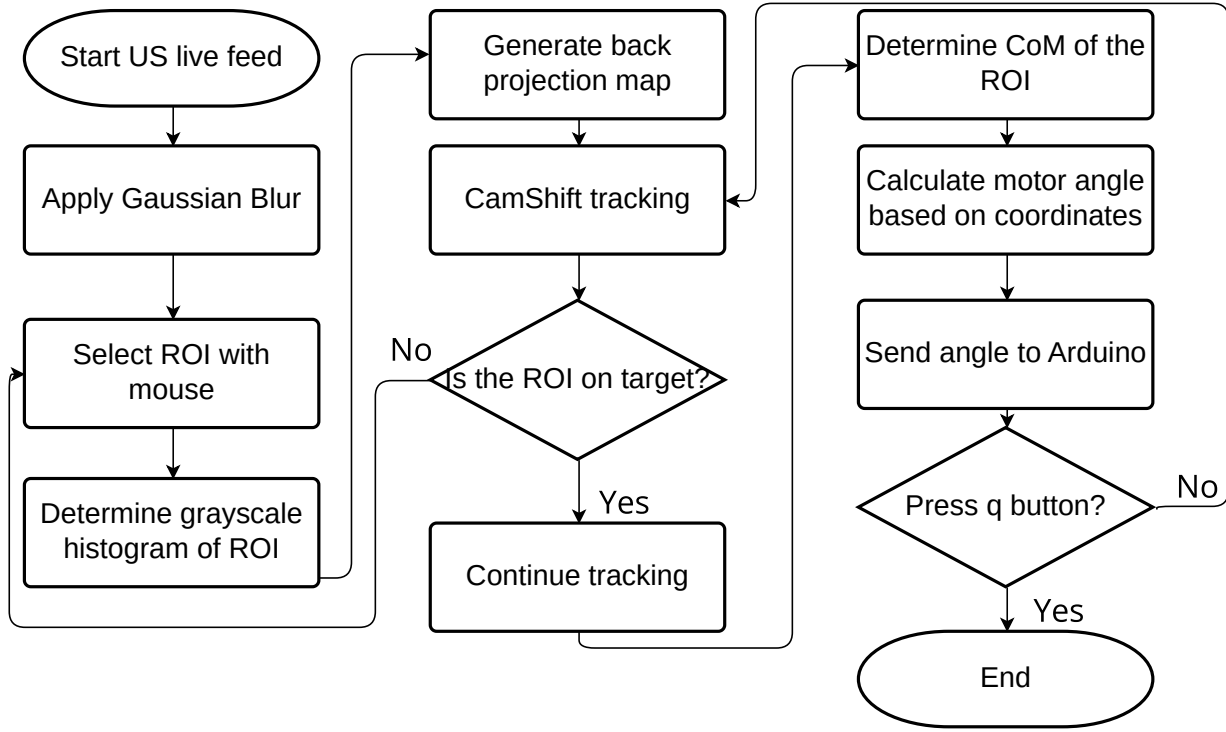


Fig. 8: A flowchart of the entire algorithm highlighting the different operations. It starts with displaying the live feed after which the tracking is performed and the motor is controlled with the pixel coordinates obtained from the tracker.

a better approximation of the CoM of the target, but could also be more vulnerable to noise.

To decrease the effect of noise or objects similar in intensity to the target, the mask is set to zero for the part of the frame 50 pixels above and 50 pixels below the initial ROI. This corresponds to almost 20% of the height of the frame since the total height is 512 pixels. This addition is made since the lesion cannot change in depth, apart from some change due to tissue deformation. This ensures that noise below and above the lesion is removed. 50 pixels (corresponding to 3.9 mm) was chosen since this will make sure the vast majority of lesions will be fully visible on the mask. Making the size a lot bigger can introduce more noise while not adding a lot of benefits (since such a large lesion should be big enough to be targeted correctly by the needle guide). Making the size smaller would also not be beneficial since it could mean that a larger portion of a big lesion is not accounted for during tracking, and it would be less forgiving on deformations.

3.2.2 Motor control

The motor, a Futaba s3003 (Futaba Corp., Mobara, Japan), is controlled with an Arduino UNO and an Adafruit motor shield V2 (Adafruit Industries, New York City, United States). The position of the target's CoM is known since the algorithm keeps track of it. However, this location is given in pixel coordinates within the transducer frame. To define the coordinates in a relevant way for the servo motor equation 1 can be used. In this case, P^t is already known. As such a definition for the transformation matrix is needed to get the coordinates in the coordinate frame of the servo. The two frames can be seen in figure 9.

The origin of the servo is used since this is irrespective of any rotation. Since the needle is positioned left of the axis in figure 9 a correction needs to be made to account for this offset.

Since the transformation from the transducer frame to the servo axis frame does not involve any rotation and only involves a displacement of the origin and a scaling factor (pixels to millimeters) the transformation matrix will look as follows:

$$\begin{bmatrix} s_x & 0 & d_x \\ 0 & s_y & d_y \\ 0 & 0 & 1 \end{bmatrix}$$

where d_x and d_y are the origin displacements in the x (width) and y (depth) direction of the origin in the servo frame relative to the origin of the transducer frame and s_x and s_y are the scaling factors in the x and y direction, respectively.

To calculate d_x and d_y it is necessary to determine the displacement from the transducer origin to the servo origin. From the CAD model, it can be derived that (x,y) equals (32.90, 11.25) mm.

To determine the scaling the scanning width and depth can be divided by the amount of pixels in the US image.

The amount of pixels can be determined with a Python function and by checking the resolution of the frame. It is 512 pixels in both directions.

The used transducer is the L12-5L40S-3 [45]. The width is 40 mm and the depth, which can be set between 20 mm to 70 mm, is set at 40 mm to create a square frame. Dividing




Fig. 9: The two frames that are used for the needle guide with the x-axis (red and horizontal) and y-axis (green and vertical). The transducer frame's origin is located on the (0,0) point of the pixels on the US feed and is used to track the target. The servo frame, placed in the center of the axis, is used to determine the necessary angle to aim the needle towards the point P.

these by the pixel values gives the following values in the matrix:

$$H_t^s = \begin{bmatrix} 0.078125 & 0 & 32.90 \\ 0 & 0.078125 & 11.25 \\ 0 & 0 & 1 \end{bmatrix} \quad (2)$$

Using this matrix and the coordinates of the point in the frame of the transducer the coordinates can be defined in the frame of the needle guide.

Once this is done the angle can be calculated from these coordinates. Since the distances in x (adjacent) and y (opposite) directions are now defined the arctangent can be used to calculate the angle.

However, this is the angle between the axis of the servo and the target. As mentioned, the needle has an offset to the left of the axis of 8.08 mm, illustrated in Figure 10.

This can be accounted for with equation 3. o_n is the needle offset in millimeters and θ_s is the servo angle in degrees. The correction is subtracted from the initially calculated angle.

$$\theta_s = \arctan\left(\frac{y_p^s}{x_p^s}\right) - \arcsin\left(\frac{o_n}{\sqrt{x_p^{s2} + y_p^{s2}}}\right) \quad (3)$$

This calculated value is then communicated to the Arduino. To reduce the possibility of errors the angle is sent in a




Fig. 10: To compensate for the offset of the needle 8.08 mm to the left, the angle (red arc) needs to be calculated. This is done with the arcsine since the opposite side (offset) and the hypotenuse (distance of the axis to the target) are known.

specific format. An angle of 30.45 degrees is formatted as "A30.45*". This angle is then mapped to a microsecond value used for the servo pulse. Using the angle directly can result in an error since every servo is slightly different. As a result, angles can deviate multiple degrees, which was the case for this servo. So, using a calibration code for the Arduino and servo one can visually check how many microseconds the pulse, supplied to the servo, should be to result in angles of 0 and 90 degrees. The other angles can then be mapped in this range.

The servo can only rotate counterclockwise from the 0-degree point and should operate from 0 to 60 degrees. So during calibration, the 0 point ensures the needle points east (into the transducer) and the 90-degree point ensures the needle points north. Before using the needle guide the calibration code can be used to visually validate that the microsecond values corresponding to 0 and 90 degrees are still correct.

To visually aid the operator during the procedure a projection is made for the needle path, represented by a semi-transparent red line based on the computed angle and point where the needle enters the frame.

The entrance point can be determined using Equation 4 with the scaling factors and origin translations from the transformation matrix.

$$y_{int} = \tan(\theta_s)(d_x + o_n \sin(\theta_s))s_x - (d_y - o_n \cos(\theta_s))s_y \quad (4)$$

y_{int} is the intersection point on the US frame in pixels. The horizontal distance between the transducer and the servo axis is known. An extra factor is necessary to compensate for the offset. In the horizontal direction, this offset is o_n when the servo angle is 90 degrees (pointing

north) and 0 when it is 0 degrees (pointing east). For the vertical direction, this is flipped. The angle is known (Equation 3), so the vertical distance to the frame's border can be determined.

The distance from the servo axis to the probe's origin is also known. After accounting for the offset this distance can be subtracted from the total vertical distance. The result needs to be scaled from millimeters to pixels. This final value, y_{int} , is the needle's entry point on the US frame.

Using this point and the known angle the path of the needle can be extrapolated and shown on the US frame.

During testing on the phantom the difference in speed of sound in the plastisol phantom compared to human tissue needs to be accounted for [46]. The scan is made with the assumption that the speed of sound in the tissue is 1540 m s^{-1} , but the speed of sound in plastisol is around 1400 m s^{-1} . When testing on a plastisol sample with a thickness of 45 mm the US image displayed a thickness of around 50 mm. So a correction needs to be made to the height coordinate used for the motor control as it needs to be scaled to compensate for this difference.

Initially, the needle did not align with the CoM and the projected path. Due to this, a calibration was proposed. Since the deviation was mostly in the y direction the y-value used for the angle calculation was multiplied with a factor to ensure the needle would hit the target. After multiple calibration attempts it was concluded that a calibration in the y-direction, used in equation 3, of 1.265 would successfully and consistently hit the target. Calibration values below 1.15 and above 1.30 would often miss the target. Values inside this range, like 1.25 and 1.28 did hit the target but during calibration, it could be seen that they slightly deviated from the predicted path. At a value of 1.265, the needle seemed to align the best with the projected path of the needle, which crosses the target's CoM.

3.3 Experimental setup

Two kinds of experiments will be performed. First, the tracker performance will be evaluated with ex vivo and in vivo tests. Second, the needle guide performance will be assessed with ex vivo tests.




Fig. 11: The performance is evaluated by tracking a tumor in a breast phantom.

3.3.1 Ex vivo

Ex vivo testing is done on a breast phantom, see Section 3.3.2 and Figure 11. This plastic phantom can contain tumors, represented by plastic objects with varying scattering levels. Since this can be created and manipulated by the maker it is easier to test on this than on a human, because the image will be less noisy and borders between tissues are better defined. Moreover, it is also more convenient than testing on a human.

First, the tracker will be tested by selecting a lesion on the live video feed. The performance can be determined by checking how quickly and accurately the tracking works. It can also be checked if the target can still be found after it exits and returns to the feed.

For the first test, a video of the tracker on different lesions is made. The probe is first held stationary. Then the probe is moved at varying speeds and distances. These tests are done to test the robustness of the tracker.

Additional tests are done where the probe is moved in a slow and controlled manner before being held still for a couple of seconds. This could give a better representation of the performance of the tracker during a normal procedure.

The position of the CoM of the ROI is obtained from the algorithm. The CoM of the lesion is approximated by drawing a rectangular ROI around the lesion. The start point and dimensions of the box are known so the CoM can be calculated. The Euclidean distance is then used to determine the deviation of the indicated CoM and the real CoM. This is done for multiple frames per video to see how the deviation changes over time. The mean and max deviation can then be obtained, as can the diagonal length of the lesion. Dividing the mean and max deviations by the diagonal length gives a relative mean and max value. Since lesions differ in size this metric can better showcase the deviation.

In addition to evaluating the accuracy of the tracker the time it takes to find the target again after an occlusion is also assessed. If the lesion leaves the scanning field, due to an incorrect angle, disengaging the probe, or moving the probe away from the lesion, the tracker loses the lesion. To ensure a low procedure time the tracker should be able to find the lesion again and find it quickly. During testing the probe is disengaged, tilted, and moved away from the lesion. After the lesion becomes visible again in the scanning field, the time it takes for the tracker to find the lesion again is measured.

The delay of the system can be evaluated by filming the US feed and the test setup simultaneously with a high-speed camera on a tripod. The needle can be moved in or out of the phantom or the probe could be placed on or disengaged from the phantom. The delay between the real action and the action on the US feed can then be determined.

The parameters needed for making the confidence map can also be tuned and tested in this way. The influence of Gaussian blur, the number of bins in the histogram, scaling of the histogram, and filtering of smaller areas with high confidence can be tested and changed for an optimal mask. This can be done by altering one of the parameters and

seeing how well the mask can distinguish the target from surrounding tissue and how well the noise is reduced.

After this, the needle guidance performance can be observed by tracking the lesion and inserting the needle. The difference between the location of the CoM and the needle can then be measured. This distance is measured perpendicular to the needle. The projected path of the needle can also be displayed as a visual aid for the operator. Additionally, biopsies can be taken of the phantom to check how often a correct sample is taken. Since the tumors are colored this could be verified visually. This gives a binary success or failure result.

3.3.2 Breast phantom

For the ex vivo testing, a breast phantom was made to test the tracking of the needle. The breast has a harder outer layer representing the skin and a softer inside. The skin is made with soft plastisol with 2 percent by mass of silica gel (60-200 μm) to create some scattering. This is heated to 160 degrees Celsius in a microwave and poured into a 3D-printed mold in a water bath. A 3D-printed breast is placed in the mold to create the skin around the breast.

The inside is made similarly to the skin but with soft plastisol with 30 percent by mass of softener added and no silica. To ensure it attaches to the skin but does not penetrate the skin, it should be added when the skin has almost cooled down. When this has cooled down to below 120 degrees Celsius the tumors can be added, to prevent the tumors from sinking a lot.

Additionally, to create a more challenging US scan, one can add ribbons of the same material as the skin to create more scattering in the breast. When this has almost cooled down another layer of the skin should be added to cover the bottom of the breast. The final result can be seen in figure 12a

Three types of tumors were made.


Blue tumors: Consist of soft plastisol heated to 160 degrees Celsius in a microwave. No silica gel was added, to prevent scattering. However, the air bubbles were not sucked out so the tumor is mostly anechoic with specks of white.

Green tumors: Created in the same way as the blue tumor. 0.4 percent by mass of silica gel was added and air bubbles were sucked out with a vacuum pump. This tumor is closer to a hypoechoic tumor.

Red tumors: Made the same way as the other tumors, but with 5 percent by mass of silica gel added. Air bubbles were removed using the vacuum pump. This tumor is hyperechoic. Due to the amount of silica, the distribution was homogenous since more silica settled on the bottom of the mold.

These tumors are added to various depths and different shapes to test the robustness of the tracking algorithm concerning the depth of the tumor, the shape of the tumor, and the grayscale intensity of the tumor.

A second phantom, seen in figure 12b, is used for additional testing since this phantom has tumors with lower contrast to the surrounding tissue. This phantom was made before this research. The same mold and approach were used as



(a) First breast phantom used for testing. (b) Second breast phantom used for testing.

with the first phantom. This phantom contains hypoechoic tumors of various degrees and shapes. Some are a lot darker and some a lot lighter. The scattering in the tissue lowers the contrast with the lighter hypoechoic tumors. A third phantom, similar to the second phantom, was used to test the accuracy of the needle guide. This phantom only contains two tumors, a small green tumor, and a blue tumor. A fourth phantom was used to calibrate the needle to align with the projected path and the CoM. This phantom, similar to the second phantom, has no tumors. The ROI was placed at a fixed point and the needle was inserted to check how well it lined up with the projected path.

3.3.3 In vivo

For in vivo testing, videos were made of the author's body. Videos were made of the carotid artery on the neck and a hematoma on the shin. Tracking these features could be a lot more challenging because, unlike the phantom, the densities, surrounding tissue, and shapes in the body cannot be controlled. In vivo and ex vivo results will be compared to see the differences in the tracker's accuracy and how quickly it finds the lesion when it re-emerges in the frame. Since it is tested on a human the needle guide will not be tested.

4 RESULTS

4.1 Tracking

The tracking performance was evaluated by selecting and tracking lesions in a phantom and on in vivo videos. Before the tracking was tested the parameters for the confidence mask had to be tuned. The experiments and results needed for the tuning are presented in Appendix 1.

The tracker performance was assessed by measuring the deviation between the real CoM of the tumor and the CoM of the ROI. Additionally, the (relative) mean and maximum deviations were calculated. The following tests are performed to test the robustness of a tracker. The probe is held stationary, moved around fast and slow, and disengaged and re-engaged to see how fast the tumor is localized again. The following phantoms and tumors were selected as targets for the tracking tests:

- 1) Test 1: phantom 1 with a hyperechoic spherical tumor, as shown in Figure 13.
- 2) Test 2: phantom 1 with a hypoechoic ellipsoidal tumor. It has quite some contrast with surrounding




Fig. 13: Ultrasound frame of the first test using Mean Shift. The red line indicates the needle path and the blue box with the cross indicates the ROI with the CoM. This tumor was also used for the CamShift tests.

tissue, but there are also similar objects in the frame. See Figure 22a.

- 3) Test 3: phantom 2 with a hypoechoic cubical tumor. This tumor is darker and larger and has significantly less contrast with surrounding tissue than the tumor in test 2. It is the largest tumor used for testing. See Figure 22b
- 4) Test 4: phantom 2 with a hyperechoic tumor. This tumor has less contrast with the surroundings than test 1, so it should be more challenging. It is also the smallest tumor that is used for testing. See Figure 22c
- 5) Test 5: phantom 1 with a cubic hypoechoic tumor. It has some contrast with the darker immediate surroundings, but less contrast with the deeper tissue and the skin. This could pose a challenge when the algorithm needs to find the tumor when it leaves and re-enters the frame. See Figure 22d

The results for Mean Shift can be seen in Tables 1 and 2. The tables show the size of the tumor, the mean deviation, and the max deviation. They also present the relative mean and max deviation to account for differences in tumor size. Each test has between ten and twenty data points.

The CamShift results can be found in Tables 3 and 4. The CamShift algorithm was selected for the needle guide tests since it has smaller deviations than the Mean Shift algorithm. Most CamShift measurements show a deviation below 2 mm, but there are outliers. These mostly occur right after the tracker finds the lesion again after an occlusion. The deviations for the CamShift algorithm across five tests are shown in Figure 14. Figure 15 shows the corresponding boxplots with the bigger outliers ($Z\text{-score} > 1$) removed. The deviation is 1.00 ± 2.26 mm (range: 0.08 mm to 15.3 mm) for the ex vivo tests and 1.36 ± 2.13 mm (range: 0.00 mm to 11.6 mm) for the in vivo tests. The big standard deviation and range are due to the few wild outliers present. Since these large deviations occur after wild movements or occlusions

we could also ignore them (using the Z-score) to get a representation that would correspond more to a normal procedure. In this case, the deviation becomes 0.55 ± 0.61 mm (range: 0.08 mm to 3.12 mm) for the ex vivo tests and 0.83 ± 0.61 mm (range: 0.00 mm to 3.12 mm) for the in vivo tests.




Fig. 14: Deviations of the CamShift algorithm for the in vivo and ex vivo tests. The in vivo tests were performed on a leg hematoma and a carotid artery. The ex vivo tests were performed on five different tumors across two breast phantoms.




Fig. 15: A boxplot representing the same data as Figure 14. Outliers ($Z\text{-score} > 1$) were removed to better visualize the median, upper quartile, and lower quartile deviations.

The performance of the Mean Shift and CamShift algorithm during a more controlled procedure was also assessed. The ex vivo test was performed on the tumor from test 4, seen in Figure 22c. The in vivo test was done on the hematoma in the leg. The mean, maximum, and minimum deviations per test can be seen in Table 5. The results are visualized in Figure 16.

This shows an increased accuracy and precision for the CamShift algorithm compared to Mean Shift. During wild movements or when the target is found right after an occlusion, the CamShift algorithm could deviate from the target more than the Mean Shift algorithm. However, during a normal, controlled procedure this does not occur and it becomes clear that CamShift is more advantageous for this application.

The time it takes for the tracker to find the lesion after occlusion was measured for both the CamShift and Mean Shift algorithms for the in vivo and ex vivo experiments. The results can be found in Figure 17. The CamShift algorithm outperformed the Mean Shift algorithm in this case.

Most of the time there are specks of noise or other objects similar to the target in the frame. In this case, the CamShift algorithm caused the ROI to become big and take in some noise and the target, before centering on the target again.

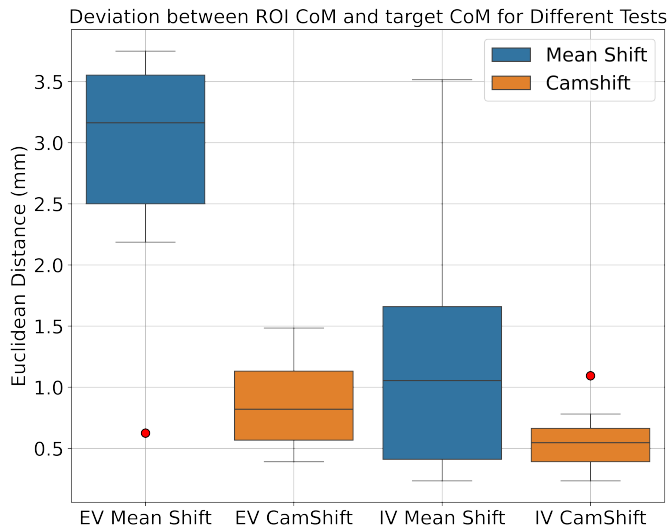


Fig. 16: Boxplots comparing the performance of CamShift (orange) and Mean Shift (blue) for both in vivo and ex vivo experiments. These tests were performed in a slow and controlled manner before keeping the probe stationary.

The Mean Shift algorithm does not have an expanding ROI. So, sometimes it got stuck in a local minimum. If this local minimum was close to the target it would shift to the target. However, if the local minimum was far away (so deeper tissue or a part of the skin) or was a comparable object bigger than the lesion, the tracker would get stuck and the target needed to be indicated again by the operator.

When the lesion did not move significantly from its last position when appearing again, so when changing the tilt or moving outside of the scanning field, the tracker had no problem finding it. For instances where the probe was disengaged and re-engaged in such a way that the location of the lesion was completely different, the tracker had more trouble finding the target. As seen in Figure 17 the CamShift algorithm found the target in under two seconds in most cases, which is an acceptable delay since it barely affects the procedure.

The framerate is around nine FPS (frames per second). However, this does not accurately reflect the total delay between an action in the scanning field and its appearance on the US frame. To determine this delay a high-speed camera was used. This showed a delay of roughly 150 ms. This delay is acceptable for the application and is similar to the 100 ms delay mentioned in Section 2.

4.2 Needle guide

Tests were done on multiple tumors in the second phantom. The tests were done on various tumors in three different phantoms. The results can be seen in Table 6.

The US live feeds of eleven of the thirteen tests were recorded. Seven of these videos could be used to determine the needle accuracy. In the remaining four videos, the tumor was either too close to the entry point of the needle or the needle was hardly/not visible. When the needle entered the tumor the path of the needle was extrapolated to determine the deviation of the needle perpendicular to the CoM of the tumor. The tests showed a deviation of 0.636 ± 0.369 mm (range: 0.228 mm to 1.18 mm).




Fig. 17: Boxplots comparing how fast the trackers find a target after occlusion. The tests were done for the CamShift (orange) and Mean Shift (blue) algorithms for both in vivo and ex vivo experiments.

5 DISCUSSION

Results show that the tracking using the CamShift algorithm results in deviations below 2 mm and that the CoM of the tracker is consistently located on the lesion when used in a controlled manner. The in vivo and ex vivo tests with wild movements and occlusions of the lesion show that the tracker is robust to quick and slow movements and can quickly find the lesion after losing sight due to occlusions. However, the tracker can have problems with large deformations and if the surrounding tissue is comparable in intensity to the target it can introduce a deviation in the CamShift algorithm. This occurs most often during quick and big movements or right after finding the target after an occlusion. In this case, surrounding noise can be seen as part of the target, resulting in an increase of the ROI and a deviation of the window from the CoM of the lesion. This problem is often resolved within a few seconds once the probe is stationary. Since the probe is stationary during the needle insertion these are only minor problems, although an effort can be made to remove the noise more efficiently. However, deformation occurs during needle insertion, making it beneficial to find a solution to make the tracker more robust to deformations.

The in vivo and ex vivo tests with normal movements of the probe show that the mean errors are 0.568 ± 0.231 mm and 0.844 ± 0.347 mm, respectively. Moreover, this is the accuracy during the entire procedure. Once the probe is stationary the error becomes smaller with the CamShift algorithm. Since tumors targeted for biopsies are often at a minimum of 5 mm in diameter, this accuracy indicates a high success rate for lesion tracking.

During the testing of the needle guide, the needle consistently entered the phantom too shallow, meaning that the angle seemed correct but the point of entry was too close to the probe. This caused the needle to miss the CoM and in some cases the lesion. This indicated a systematic error in the system. More research is required to identify the source of this error.

It should be noted that the servo calibration was done visually, introducing a human error since one cannot accurately see if the servo angle is at 0 and 90 degrees during the calibration. An alternative can be proposed where the calibration is done with a camera and a marker or with a laser and a reflective surface on the needle holder. This should improve the servo calibration

Table 6 shows that most biopsies performed on different lesions, result in specimens that contain parts of the lesion. Two of the tests resulted in a failed biopsy without any tissue. However, on the US feed, the needle could be seen penetrating the tumor. The lesion in this case was small and perforated multiple times before. Whether the system caused the failure, the lack of tissue in the tumor, or the operator's inexperience in using the needle mechanism, was not clear in this case.

The tests to evaluate the system's delay were done with a high-speed camera. However, with the setup it was hard to get great visuals of the screen and the exact moment the needle was inserted or when the probe was placed on the phantom. It was also hard to link this exact moment to the same moment on the screen. When the probe is disengaged it is not removed instantaneously. There is a gradual (but quick) movement where the pressure of the probe on the phantom becomes smaller and smaller. The angle needed to film both this movement and the screen makes it possible to see when the probe is being disengaged but it is hard to see the exact moment it disengages. Similarly, on the screen one can see the moment the probe is firmly pressed on the phantom and the moment it is removed, but the time in between also takes 20 ms to 30 ms. The measured delay seems to be around 150 ms but could have a deviation of 20 ms to 30 ms. Using a setup that shows the screen and the probe better or using a square phantom with fewer impurities could improve the accuracy of the delay measurements.

When losing sight of the lesion the tracker sometimes needs a couple of seconds to find the lesion again. Improving the tracker by using more information, like geometry or edges, could ensure that fewer parts in the frame correspond to the target. This could result in a faster localization of the target. An attempt was made to add geometric information by combining the intensity histogram with a histogram of oriented gradients (HOG). However, this caused the frame rate to drop to two frames per second, which is insufficient for this application. Improving the frame rate could also increase the speed of the algorithm. Since the algorithm was written in Python switching to C++ or C could be beneficial to reduce execution time and increase the framerate.

The needle path projection is a helpful visual aid for the operator. Two shortcomings are that the needle could deviate from the path and that it is impossible to determine the distance between the needle tip and the target's CoM since it only indicates the predicted path. Using needle detection could be an alternative to determine the distance and visualize the needle in real-time. In this case, an extra motor could be added to automate the insertion. The downside is that this would be more computationally expensive and should be very robust and precise.

During biopsies of a lesion, the ROI might start moving once the needle is inserted in the tumor, due to deformations. The algorithm can be stopped as soon as a sample has been taken, to prevent the motor from moving when retracting the needle. An addition could be made where the motor angle gets locked during insertion so it cannot move in case the tracker starts to deviate due to deformations

The accuracy of this research is comparable to the accuracy in cited studies (1.1 mm to 2.9 mm for in vivo tests and 1.8 mm for ex vivo tests) [19], [22], [25], [26]. The accuracy could be increased by improving the system calibration and by improving the tracker to be more robust to deformities during the needle insertion.

It should be noted that the needle accuracy is only tested on a breast phantom. The accuracy values from other US-assisted robotized needle guides were obtained from in vivo tests. The needle accuracy for in vivo tests is expected to be lower than for ex vivo tests.

A real breast is more deformable than the phantoms used for this research. It is possible to make a more accurate phantom by making it softer. However, this would mean that the phantom is less stable and cannot be punctured as often as a more rigid breast. So fewer tests can be done on one breast, making this process more labor intensive since more phantoms need to be made. Additionally testing the initial performance of the tracker and needle guide would have been harder on a very soft phantom. For these reasons, only stiffer phantoms were used during this research.

The main reason this stiffness is relevant is because one problem that was observed was the deformation of the tissue during the needle insertion. This caused the lesion to deform and the tracker could start to deviate from the lesion. Because a real breast is more deformable it is expected that the tracker will have more trouble staying on the lesion during needle insertion. This can negatively impact the accuracy.

6 CONCLUSION

This research offers an accurate and simple US-assisted robotized needle guide for breast biopsies. The handheld device can assist an operator by pointing the needle toward the lesion. This can reduce room for error and operating time.

The developed system utilizes the CamShift algorithm with back projection to track a user-defined rectangular ROI. The servo motor is controlled using an Arduino Uno, and inverse kinematics and trigonometry are applied to point the needle to the lesion to assist an operator. The application is brought together in a functional design, designed in SolidWorks.

The CamShift tracking showed an error sub 2 mm when using it in a controlled manner for both in vivo and ex vivo tests. To test the robustness the probe was moved in a less controlled manner. In this case, the probe was moved with varying speeds and jerky movements, and the target was moved outside the frame. For this approach the mean error was also below 2 mm, although there were more outliers. Since biopsies are mostly performed on lesions of 5 mm and above this result means that the tracker performs well enough to target such a lesion.

Additionally, the algorithm quickly found the target again after occlusion (below 2 seconds most of the time). With a framerate of roughly nine FPS and a total delay of around 150 ms, the device is fast enough to not hinder an operator during the procedure. The tracking algorithm performed, in most cases, better for the ex vivo tests. This can be attributed to the bigger contrasts and the smaller deformations in the phantom compared to human tissue.

The needle guide, with a needle accuracy of 0.6 ± 0.4 mm, is comparable to the state of the art. It should be noted that most of the state-of-the-art results are based on in vivo biopsies. The needle accuracy results in this study are based on ex vivo tests. Since a breast is more deformable than the used phantoms, it can be expected that the accuracy will decrease when tested in vivo.

This system is an attractive option compared to state-of-the-art robots for various reasons. It is handheld, cheap, and simple, it uses US which is noninvasive, real-time, quick, and does not expose the patient to ionizing radiation. A disadvantage is that certain lesions, that might be visible on MRI or CT scans, could be invisible on US scans. Additionally, the device is handheld, and needle insertion is performed by an operator. Because of this, there is still room for human error, which depends on the operator's skill.

Overall the presented device is a real-time non-invasive and accurate option for breast biopsies.

6.1 Future research

To improve the accuracy of the device, a couple of additions and changes could be made to the tracker and the needle guide.

Currently, the tracker solely works on information regarding the intensity of the target. Other parts in the frame can have a similar intensity, which would result in unwanted noise. To prevent this, geometric information like edges or corners could be used to represent the target more accurately. This could improve the tracking of the target. One foreseeable issue is that, during needle insertion, the lesion deforms. If the algorithm also takes shape/geometry into account, this could negatively impact the tracking. So if this were to be applied it would be necessary to update the geometric information to account for deformations continuously. On top of that, segmentation or machine learning could be employed to assist the operator with automatic lesion or needle detection.

Furthermore, the amount of human error could be reduced by adding an extra motor to insert the needle. However, this would increase the complexity of the system. Additionally, the calibration method of the servo could be improved to ensure that the insertion angle is correct.

To improve the needle accuracy the servo calibration could be done with a camera and a marker to decrease the error in the servo angle.

With these possible improvements incorporated the system could have a positive effect in clinical use by reducing operation time and cost and by increasing the accuracy of the biopsy.

REFERENCES

- [1] I. Jung, K. Han, M. J. Kim, H. J. Moon, J. H. Yoon, V. Y. Park, and E.-K. Kim, "Annual trends in ultrasonography-guided 14-gauge core needle biopsy for breast lesions," *Korean journal of radiology*, vol. 21, no. 3, p. 259, 2020.
- [2] G. M. Fuhrman, G. J. Cederbom, J. S. Bolton, T. A. King, J. L. Duncan, J. L. Champaign, D. H. Smetherman, G. H. Farr, R. R. Kuske, and W. M. McKinnon, "Image-guided core-needle breast biopsy is an accurate technique to evaluate patients with nonpalpable imaging abnormalities," *Annals of surgery*, vol. 227, no. 6, pp. 932–939, 1998.
- [3] M. F. Dillon, A. D. Hill, C. M. Quinn, A. O'Doherty, E. W. McDermott, and N. O'Higgins, "The accuracy of ultrasound, stereotactic, and clinical core biopsies in the diagnosis of breast cancer, with an analysis of false-negative cases," *Annals of surgery*, vol. 242, no. 5, pp. 701–707, 2005.
- [4] L. W. Bassett, M. C. Mahoney, and S. K. Apple, "Interventional breast imaging: current procedures and assessing for concordance with pathology," *Radiologic Clinics of North America*, vol. 45, no. 5, pp. 881–894, 2007.
- [5] J. H. Youk, E.-K. Kim, M. J. Kim, and K. K. Oh, "Sonographically guided 14-gauge core needle biopsy of breast masses: a review of 2,420 cases with long-term follow-up," *American journal of roentgenology*, vol. 190, no. 1, pp. 202–207, 2008.
- [6] J. E. Masciopinto, A. B. Levin, M. P. Mehta, and B. S. Rhode, "Stereotactic radiosurgery for glioblastoma: a final report of 31 patients," *Journal of neurosurgery*, vol. 82, no. 4, pp. 530–535, 1995.
- [7] "American cancer society," 2022.
- [8] V. Groenhuis, *Robotic systems for breast biopsy using MRI and ultrasound imaging*. 2020.
- [9] F. J. Siepel, B. Maris, M. K. Welleweerd, V. Groenhuis, P. Fiorini, and S. Stramigioli, "Needle and Biopsy Robots: a Review," *Current Robotics Reports*, vol. 2, no. 1, pp. 73–84, 2021.
- [10] K. Tamaki, H. Sasano, T. Ishida, M. Miyashita, M. Takeda, M. Amari, N. Tamaki, and N. Ohuchi, "Comparison of core needle biopsy (cnb) and surgical specimens for accurate preoperative evaluation of er, pgr and her2 status of breast cancer patients," *Cancer science*, vol. 101, no. 9, pp. 2074–2079, 2010.
- [11] G. Schueller, C. Schueller-Weidekamm, and T. Helbich, "Accuracy of ultrasound-guided, large-core needle breast biopsy," *European radiology*, vol. 18, pp. 1761–1773, 2008.
- [12] E. E. Deurloo, K. G. Gilhuijs, L. J. S. Kool, and S. H. Muller, "Displacement of breast tissue and needle deviations during stereotactic procedures," *Investigative radiology*, vol. 36, no. 6, pp. 347–353, 2001.
- [13] A. Lakoma and E. S. Kim, "Minimally invasive surgical management of benign breast lesions," *Gland surgery*, vol. 3, no. 2, p. 142, 2014.
- [14] L. J. Brattain, C. Floryan, O. P. Hauser, M. Nguyen, R. J. Yong, S. B. Kesner, S. B. Corn, and C. J. Walsh, "Simple and effective ultrasound needle guidance system," in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 8090–8093, IEEE, 2011.
- [15] L. Tielens, R. Damen, J. Lerou, G.-J. Scheffer, and J. Bruhn, "Ultrasound-guided needle handling using a guidance positioning system in a phantom," *Anaesthesia*, vol. 69, no. 1, pp. 24–31, 2014.
- [16] G. J. Vrooijink, M. Abayazid, and S. Misra, "Real-time three-dimensional flexible needle tracking using two-dimensional ultrasound," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 1688–1693, 2013.
- [17] C. K. Min Li and M. Stone, "Automatic contour tracking in ultrasound images," *Clinical Linguistics & Phonetics*, vol. 19, no. 6-7, pp. 545–554, 2005. PMID: 16206482.
- [18] B. M. Maris, P. Fiorini, A. Calanca, C. Tenga, F. J. Siepel, V. Groenhuis, and S. Stramigioli, "Increasing the precision of the biopsy with robots: two case studies," in *2nd Italian Conference on Robotics and Intelligent Machines, I-RIM 2020*, pp. 1–4, 2020.
- [19] M. Welleweerd, D. Pantelis, A. de Groot, F. Siepel, and S. Stramigioli, "Robot-assisted ultrasound-guided biopsy on mr-detected breast lesions," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2965–2971, 2020.
- [20] G. Kronreif, M. Fürst, J. Kettenbach, M. Figl, and R. Hanel, "Robotic guidance for percutaneous interventions," *Advanced Robotics*, vol. 17, no. 6, pp. 541–560, 2003.

- [21] J. Kettenbach and G. Kronreif, "Robotic systems for percutaneous needle-guided interventions," *Minimally Invasive Therapy & Allied Technologies*, vol. 24, no. 1, pp. 45–53, 2015.
- [22] J. Kettenbach, G. Kronreif, M. Figl, M. Fürst, W. Birkfellner, R. Hanel, and H. Bergmann, "Robot-assisted biopsy using ultrasound guidance: initial results from in vitro tests," *European radiology*, vol. 15, pp. 765–771, 2005.
- [23] S. Jacobo-Zavaleta and J. Zavaleta, "Needle placement for robot-assisted 3d-guided ultrasound breast biopsy: A preliminary study," *IEEE Latin America Transactions*, vol. 21, no. 3, pp. 450–456, 2023.
- [24] K. Liang, A. J. Rogers, E. D. Light, D. von Allmen, and S. W. Smith, "Three-dimensional ultrasound guidance of autonomous robotic breast biopsy: feasibility study," *Ultrasound in medicine & biology*, vol. 36, no. 1, pp. 173–177, 2010.
- [25] G. Megali, O. Tonet, C. Stefanini, M. Boccadoro, V. Papaspyropoulos, L. Angelini, and P. Dario, "A computer-assisted robotic ultrasound-guided biopsy system for video-assisted surgery," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2001: 4th International Conference Utrecht, The Netherlands, October 14–17, 2001 Proceedings 4*, pp. 343–350, Springer, 2001.
- [26] L. Yang, S. Duan, Y. Zhang, L. Hao, S. Wang, Z. Zou, Y. Hu, S. Chen, Y. Hu, and L. Zhang, "Feasibility and safety of percutaneous puncture guided by a 5g-based telerobotic ultrasound system: an experimental study," *CardioVascular and Interventional Radiology*, pp. 1–8, 2024.
- [27] T. G. Leighton, "What is ultrasound?," *Progress in biophysics and molecular biology*, vol. 93, no. 1-3, pp. 3–83, 2007.
- [28] D. Dalecki, "Mechanical bioeffects of ultrasound," *Annu. Rev. Biomed. Eng.*, vol. 6, pp. 229–248, 2004.
- [29] D. W. Townsend, J. P. Carney, J. T. Yap, and N. C. Hall, "Pet/ct today and tomorrow," *Journal of Nuclear Medicine*, vol. 45, no. 1 suppl, pp. 45–14S, 2004.
- [30] C.-Y. Liang, T.-B. Chen, N.-H. Lu, Y.-C. Shen, K.-Y. Liu, S.-Y. Hsu, C.-J. Tsai, Y.-M. Wang, C.-I. Chen, W.-C. Du, and Y.-H. Huang, "Classification of breast ultrasound tomography by using textural analysis," *Iranian Journal of Radiology*, vol. In Press, 04 2020.
- [31] A. Manbachi and R. S. Cobbold, "Development and application of piezoelectric materials for ultrasound generation and detection," *Ultrasound*, vol. 19, no. 4, pp. 187–196, 2011.
- [32] M. M. Quien and M. Saric, "Ultrasound imaging artifacts: How to recognize them and how to avoid them," *Echocardiography*, vol. 35, no. 9, pp. 1388–1401, 2018.
- [33] M. M. Aracava, R. Chojniak, J. A. Souza, A. G. Bitencourt, and E. F. Marques, "Identification of occult breast lesions detected by magnetic resonance imaging with targeted ultrasound: a prospective study," *European journal of radiology*, vol. 83, no. 3, pp. 516–519, 2014.
- [34] R. J. Hooley, L. M. Scoutt, and L. E. Philpotts, "Breast ultrasonography: state of the art," *Radiology*, vol. 268, no. 3, pp. 642–659, 2013.
- [35] L. W. Schelke, D. Cassuto, P. Velthuis, and X. Wortsman, "Nomenclature proposal for the sonographic description and reporting of soft tissue fillers," *Journal of cosmetic dermatology*, vol. 19, no. 2, pp. 282–288, 2020.
- [36] K. Martin, "Introduction to b-mode imaging," in *Diagnostic Ultrasound, Third Edition*, pp. 1–5, CRC Press, 2019.
- [37] D. Comaniciu, V. Ramesh, and P. Meer, "Real-time tracking of non-rigid objects using mean shift," in *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, vol. 2, pp. 142–149, IEEE, 2000.
- [38] B.-C. Seo, S.-S. Kim, and D.-Y. Lee, "Target-tracking system for mobile surveillance robot using camshift image processing technique," *Transactions of the Korean Society of Mechanical Engineers A*, vol. 38, 02 2014.
- [39] S. M. Nadgeri, S. Sawarkar, and A. D. Gawande, "Hand gesture recognition using camshift algorithm," in *2010 3rd International Conference on Emerging Trends in Engineering and Technology*, pp. 37–41, IEEE, 2010.
- [40] O.-D. Nouar, G. Ali, and C. Raphael, "Improved object tracking with camshift algorithm," in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 2, pp. II–II, IEEE, 2006.
- [41] T. Wen, C. Wang, Y. Zhang, and S. Zhou, "A novel ultrasound probe spatial calibration method using a combined phantom and stylus," *Ultrasound in Medicine & Biology*, vol. 46, no. 8, pp. 2079–2089, 2020.
- [42] P.-Y. Hsiao, S.-S. Chou, and F.-C. Huang, "Generic 2-d gaussian smoothing filter for noisy image processing," in *TENCON 2007 - 2007 IEEE Region 10 Conference*, pp. 1–4, 2007.
- [43] M. Ries, B. D. De Senneville, S. Roujol, Y. Berber, B. Quesson, and C. Moonen, "Real-time 3d target tracking in mri guided focused ultrasound ablations in moving tissues," *Magnetic Resonance in Medicine*, vol. 64, no. 6, pp. 1704–1712, 2010.
- [44] A. Sobral and A. Vacavant, "A comprehensive review of background subtraction algorithms evaluated with synthetic and real videos," *Computer Vision and Image Understanding*, vol. 122, pp. 4–21, 2014.
- [45] TELEMED, "Micrus scanner," 2024.
- [46] G. M. Spirou, A. A. Oraevsky, I. A. Vitkin, and W. M. Whelan, "Optical and acoustic properties at 1064 nm of polyvinyl chloride-plastisol for use as a tissue phantom in biomedical optoacoustics," *Physics in Medicine & Biology*, vol. 50, no. 14, p. N141, 2005.


APPENDIX 1: CONFIDENCE MASK

The confidence mask shows the likelihood that a certain pixel belongs to the ROI. Certain parameters influence this mask: the number of bins for the grayscale histogram, the Gaussian blur, and the area under the histogram.

These parameters were tested by keeping the other parameters the same and varying one of the parameters to see the influence on the mask. For this experiment, the same anechoic tumor was used.


The first test was performed by varying the number of bins. The grayscale has 256 values from 0 (black) to 255 (white). Using 256 bins means that each grayscale value has its bin. Using for example 64 bins means that each bin contains four pixel values. The results can be seen in Figure 18. It can be observed that a higher amount of bins results in more detail and accuracy since only the pixels present in the ROI are selected. The reason that the mask with more bins is less visible is that the area under the curve for the histogram is the same. So more bins result in lower values per bin. This can be solved by scaling the area under the curve to a higher value

The downside to having more bins is that, due to the amount of detail, parts that belong to the tumor but are not as represented by the histogram will not be marked as high confidence. Resulting in a target that is not very homogeneous. Due to this, the tracker could shift more since the target contains holes. Fewer bins result in more false positives but also create a more homogeneous target. A balance is needed so the target has few holes and other parts are not marked as high confidence.



(a) Mask with 16 histogram bins. (c) Mask with 128 histogram bins.
(b) Mask with 64 histogram bins. (d) Mask with 256 histogram bins.
Fig. 18: Influence of number of histogram bins on the confidence mask.

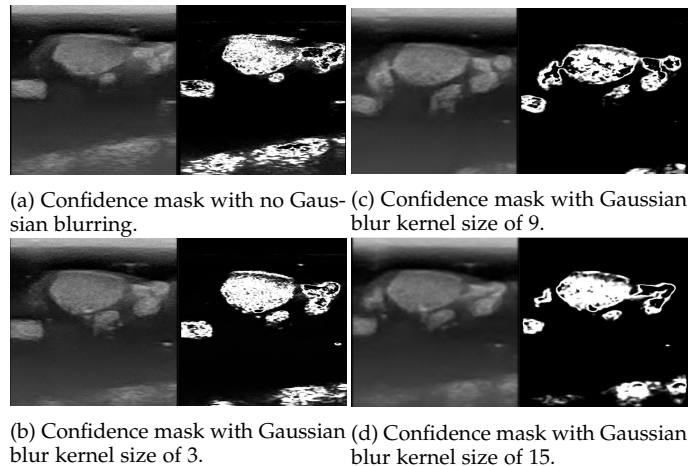
The effect of the area under the curve and its scaling can also be observed, as seen in Figure 19. The default area under the curve equals 255. This can then be scaled to make the confidence peaks more pronounced. Figure 19 shows that the default scaling was inadequate since the target was barely distinguishable from the background. Increasing the scaling gave better results. However, increasing the area under the curve too much results in a confidence map that highlights too many false positives.



(a) Confidence mask with a scaling factor of 1. (c) Confidence mask with a scaling factor of 100.
(b) Confidence mask with a scaling factor of 10 (d) Confidence mask with a scaling factor of 10000.
Fig. 19: Influence of area under the curve on the confidence mask.

The effect of Gaussian blur is shown in Figure 20. For this, the kernel size was varied while the standard deviation was kept at 0 in Python (meaning that it is automatically calculated from the kernel size).

Having no Gaussian blur results in more holes in the target. On the other hand, having a high amount of blur will reduce noise and give a more homogeneous target, but can result in more false positives or enlarge wrong areas since grayscale values will be more connected due to the blurring.



(a) Confidence mask with no Gaussian blurring. (c) Confidence mask with Gaussian blur kernel size of 9.
(b) Confidence mask with Gaussian blur kernel size of 3. (d) Confidence mask with Gaussian blur kernel size of 15.
Fig. 20: Influence of Gaussian blur on the confidence mask.

Using the results a decision was made for the various parameters. For the Mean Shift algorithm, an approach using 64 bins with a scaling factor of 10 for the area and a Gaussian blur kernel size of (3,3) was chosen since this gave a good balance between robustness and precision. This version was tested on a lesion with more challenging surroundings and less scattering. This result can be seen in Figure 21. As can be seen, it still picks up quite some noise but there is one distinct bigger confidence area. Additionally, the skin area's confidence is set to zero so the algorithm will not consider the skin a target.

When moving the probe at a normal pace, the tracker works well. If the lesion disappears from view, by removing the probe or going away from the lesion, it can find the

lesion again. When appearing again close to the last seen spot it will quickly find the target. When appearing on a completely different spot it takes a few seconds to detect the target again due to the noise and the nature of the Mean Shift algorithm.

When using the CamShift algorithm the amount of bins was increased to 256. Since this also results in less visible confidence peaks the scaling was increased to 100. The reason for this bin change is that it is more precise. The downside, as mentioned, is that this can create holes in the target when pixels in the target are missing in the histogram. For the Mean Shift algorithm, this can cause the ROI to shift within the target. Since CamShift does not suffer as much from this problem, the amount of bins can be increased.




Fig. 21: The confidence mask of a more challenging target.

This result could be further enhanced by applying a kernel and filter that connect areas within proximity and remove areas under a certain threshold. This addition was not fine-tuned, since it can be different for different tumors, tissues, and US settings. The area threshold was set to 10 pixels squared and the kernel size was set to (1,1). This is mainly done to remove specks of noise/high confidence points.

APPENDIX 2: RESULTS




Fig. 22: Ultrasound frame for the tracking tests. The blue box is the ROI where the cross indicates the CoM. The red line indicates the needle path. These tumors are used for the Mean Shift and CamShift tests.

Test	size [mm]	Mean deviation [mm]	Relative mean deviation [%]	Max deviation [mm]	Relative max deviation [%]
1	14.1, 9.96	1.60 ± 1.14	9.28 ± 6.61	3.83	22.2
2	15.0, 7.81	1.38 ± 0.841	8.19 ± 4.97	3.36	19.9
3	17.6, 12.7	2.68 ± 1.78	12.4 ± 8.21	6.09	28.1
4	14.5, 9.02	1.11 ± 0.505	6.48 ± 2.96	1.95	11.4
5	9.92, 13.3	1.75 ± 1.38	10.5 ± 8.32	4.69	28.2

TABLE 1: The (relative) maximum and mean deviation of the tracker for different tumors.

Test	size [mm]	Mean deviation [mm]	Relative mean deviation [%]	Max deviation [mm]	Relative max deviation [%]
1	11.0, 8.82	1.89 ±1.00	13.4 ±7.08	4.45	31.5
2	12.2, 7.70	2.10 ±1.15	14.6 ±7.99	4.06	28.2
3	10.5, 8.55	1.94 ±0.954	14.3 ±7.06	3.98	29.5
4	9.80, 6.01	1.75 ±1.49	15.2 ±12.9	4.84	42.1
5	13.4, 5.59	2.62 ±1.26	18.0 ±8.72	4.92	34.0

TABLE 2: The (relative) maximum and mean deviation of the tracker for different in vivo targets.

Test	size [mm]	Mean deviation [mm]	Relative mean deviation [%]	Max deviation [mm]	Relative max deviation [%]
1	12.2, 8.55	0.195 ±0.0597	1.31 ±0.400	0.313	2.09
2	15.7, 7.38	0.263 ±0.201	1.52 ±1.16	0.781	4.51
3	16.5, 12.9	2.16 ±3.48	10.4 ±16.6	11.6	55.3
4	17.1, 8.16	0.708 ±0.460	3.73 ±2.42	1.95	10.3
5	15.5, 7.58	0.355 ±0.174	2.05 ±1.01	0.781	4.52

TABLE 3: The (relative) maximum and mean deviation of the CamShift tracker for different tumors.

Test	size [mm]	Mean deviation [mm]	Relative mean deviation [%]	Max deviation [mm]	Relative max deviation [%]
1	9.69, 6.56	1.95 ±2.66	16.7 ±22.7	10.7	91.5
2	6.75, 6.99	2.01 ±1.54	20.8 ±15.9	6.88	70.7
3	8.79, 7.15	2.31 ±3.40	20.4 ±30.0	15.3	135
4	9.80, 4.73	0.344 ±0.161	3.16 ±1.48	0.625	5.74
5	11.4, 5.55	0.547 ±0.660	4.32 ±5.22	2.03	16.1

TABLE 4: The (relative) maximum and mean deviation of the CamShift tracker for different in vivo targets.

Test	size [mm]	Mean deviation [mm]	Max deviation [mm]	Minimum deviation [mm]
Ex vivo Mean shift	14.7, 9.4	2.87 ±0.914	3.75	0.625
Ex vivo CamShift	14.8, 9.77	0.844 ±0.347	1.48	0.391
In vivo Mean shift	8.0, 4.0	1.35 ±1.16	3.52	0.234
In vivo CamShift	10.3, 4.88	0.568 ±0.231	1.09	0.234

TABLE 5: The (relative) maximum and mean deviation of the two trackers when moving the probe around in a slow and controlled manner.

Phantom	Tumor	Result
Phantom 1	light Hypoechoic	Success
Phantom 1	Anechoic, hardly visible	Success
Phantom 1	Anechoic, hardly visible	Success
Phantom 1	Anechoic (second one), hardly visible	Success
Phantom 2	Dark Hypoechoic, superficial	Success
Phantom 2	Dark Hypoechoic, deep	Success
Phantom 2	Dark Hypoechoic, deep	Success
Phantom 3	Dark Hypoechoic, deep	Success
Phantom 3	Dark Hypoechoic, deep	Success
Phantom 3	light Hypoechoic, superficial	Success
Phantom 3	light Hypoechoic, superficial	Failure
Phantom 3	light Hypoechoic, superficial	Failure
Phantom 3	light Hypoechoic, superficial	Success

TABLE 6: Biopsy results of the needle guide on various tumors.

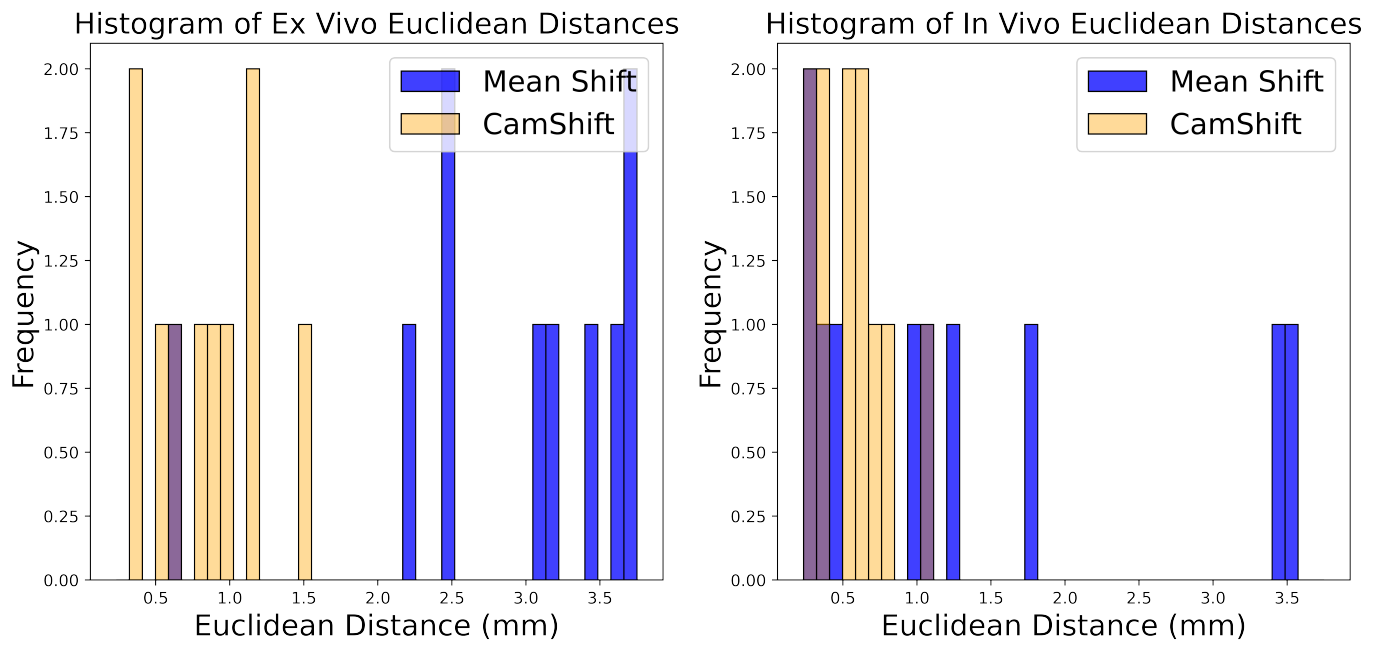


Fig. 23: Histograms comparing the performance of CamShift (orange) and Mean Shift (blue) for both in vivo and ex vivo experiments. These tests were performed in a slow and controlled manner before keeping the probe stationary.