

UNIVERSITY COLLEGE LONDON

ULTRASOUND-BASED SKULL REGISTRATION FOR TRANSCRANIAL
ULTRASOUND STIMULATION

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF MASTER OF SCIENCE (MSC)
IN PHYSICS AND ENGINEERING IN MEDICINE

BY

SIERRA M BONILLA

SUPERVISORS:

PROF. BRADLEY TREEBY (B.TREEBY@UCL.AC.UK),
DR. ANTONIO STANZIOLA (A.STANZIOLA@UCL.AC.UK)

WORD COUNT: 9569

DATE OF SUBMISSION: SEPTEMBER 2022

I, Sierra M Bonilla confirm that the work presented in this report is my own.

Where information has been derived from other sources,

I confirm that this has been indicated in the report.

Copyright © 2022 by Sierra M Bonilla

All Rights Reserved

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iv
ABSTRACT	v
1 INTRODUCTION	1
1.1 Transcranial Ultrasound Stimulation	1
1.2 Preliminary Work by Biomedical Ultrasound Group	3
1.3 Problem Statement	4
1.4 Related Works	6
1.5 Objectives	9
2 REGISTRATION ALGORITHM DEVELOPMENT	10
2.1 Ultrasound Point Cloud Reconstruction	10
2.2 Acquisition Scripts	11
2.3 Registration Method	14
3 NUMERICAL EXPERIMENTS	20
3.1 Methods	20
3.2 Number of Point Pairs	22
3.3 Noise	23
3.4 Uniform Scaling	24
3.5 Rotation	27
3.6 Translation	29
3.7 Summary of Results	30
4 EXPERIMENTAL VALIDATION USING 3D-PRINTED PVA PHANTOM	32
4.1 Methods	32
4.2 Results	36
5 PRELIMINARY EXPERIMENTAL VALIDATION USING HUMAN VOLUNTEER	39
5.1 Methods	39
5.2 Results	40
6 CONCLUSION	42
6.1 Discussion	42
6.2 Future Work	43
REFERENCES	47

ACKNOWLEDGMENTS

I would like to thank my parents for their unconditional love and support throughout my life and their hard work that made everything in my life possible. I also want to thank my grandparents for their hard work and dedication to making a better life for the generations below. A special thanks to all the women who have come before me and pathed the way for women like me to choose the paths we want, regardless of what field it may be in.

Additionally, I would like to sincerely thank both Prof. Bradley Treeby and Dr. Antonio Stanziola for their support and direction through my project. Big thanks to Prof. Bradley Treeby for making me feel like a fundamental part of the neurostim project and including me wherever possible. I'd also like to thank Dr. Antonio Stanziola for patiently answering my numerous questions regardless of the time or topic. And, I likely introduced significantly less bugs into my code thanks to ending early from Dr. Stanziola's polite suggestions to take breaks.

ABSTRACT

Purpose: Transcranial ultrasound stimulation is being explored for neuromodulation using a hemi-spherical phased element array in the Biomedical Ultrasound Group at University College London. However, the skull distorts the acoustic field. This dissertation explores the question: can an ultrasound-based image of the skull be registered to a pre-therapeutic MR or CT image to correct for the variation in the acoustic field produced by the skull?

Method: (1) Algorithms were developed and implemented to generate an ultrasound-derived point cloud of the object outer surface from the measured data. (2) Developed acquisition scripts and gathered ultrasound measurements from skull models and three subjects using a transcranial ultrasound array. (3) Determined the transformation required to co-align the ultrasound-derived point cloud with a mesh-derived point cloud from either a CAD file of the model or an MR image for human experiments.

Results: The acquisition, point cloud reconstruction, and image registration scripts were created, tested, and validated. The technique was validated using an MR image as a ground truth for location. The technique locates the object in the TUS hemi-spherical array with an error of $\pm 2^\circ$ and ± 2.9 mm from the true location. The registration algorithm was robustly tested using several experimental datasets to determine conditions under which it works best, using both experimental and simulated data. The algorithm was found to be robust to changes in the number of point pairs used, noise, and translational changes to the initial guess. However, it failed or performed worse during increased scaling and increased rotational movement experiments.

CHAPTER 1

INTRODUCTION

1.1 Transcranial Ultrasound Stimulation

Applications of neurostimulation have been of significant interest to physicians for nearly 2000 years, dating back to Scribonius Largus in 47 AD who documented the effect of transcutaneous electrostimulation from electric eels on pain modulation [1]. The purposeful modulation of the nervous system has been explored for both therapeutic purposes, examples of which include pain control, incontinence, and epilepsy; and research purposes like neural mapping. Neural activity can be modulated by invasively or non-invasively applying an electrical current or a changing magnetic field near the neural fibers. More recently focused ultrasound has proven capable of stimulating or suppressing neuronal activity non-invasively in both peripheral and central nervous system applications. Because of its preferential non-invasive quality, high spatial resolution, and ability to reach deep brain structures, transcranial ultrasound stimulation (TUS) has garnered increasing interest from the research community. An example of a single element TUS device can be seen in Figure 1.

However, transcranial ultrasound was considered impossible for many years due to the complications presented by the skull. Then in 1998, Hynynen and Jolesz showed that it is possible to correct the skull-induced aberrations to the acoustic field in *in-vitro* phased array experiments using feedback from a hydrophone placed at the focus [2]. Further work by Clement and Hynynen in 2002 showed that the distortions could be corrected for using phase shifts calculated by a wavevector-frequency domain model with acoustic properties of the skull from computed tomography (CT) scans and an accurate and fixed registration of the skull and transducer locations [3]. Since then, CT-based derivations of phase and amplitude corrections have become the gold standard, opening the door for transcranial ultrasound applications [4].

Transcranial ultrasound stimulation involves sending focused or unfocused low-intensity acoustic waves into the brain. Through mechanisms that are not entirely understood, action potentials becomes more or less likely depending on the ultrasound parameters. It is currently hypothesized that the TUS waves mechanically alter neuronal ion channel gating resulting in neuronal stimulation or suppression as seen in the Figure 2 [5].

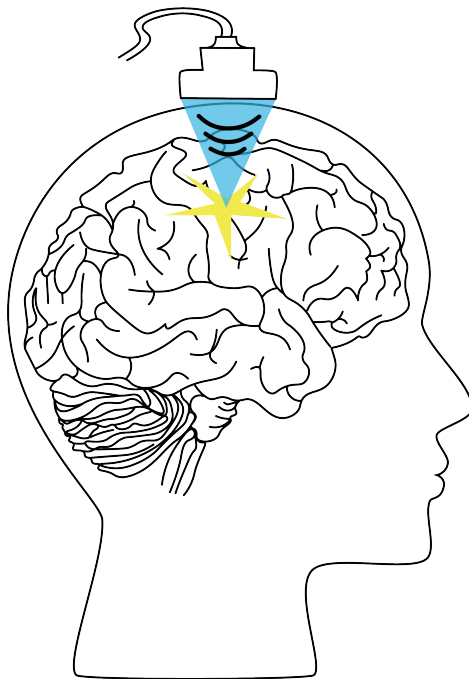


Figure 1: Single element focused TUS device used for neurostimulation.

To monitor TUS real-time functional MRI (fMRI) is regularly integrated. Unlike other stimulation techniques, TUS can target both superficial and deep regions of the brain with high spatial resolution, making applications of this technique very widespread. TUS, combined with fMRI, is especially useful for non-invasive functional neural mapping [6].

The first TUS human studies were performed by Hameroff et al. and Legon et al. in 2013 and 2014, respectively, showing transcranial ultrasound stimulation modulating cortical function, suggesting its usage as a therapy for regulating mental states or neurological disorders




Figure 2: The mechanism of neuromodulation illustrated: ultrasound (left) stimulating action potential and (right) suppressing action potential.

non-invasively [7],[8]. Transcranial focused ultrasound in conjunction with microbubbles has also been shown to facilitate targeted drug delivery by opening the blood-brain barrier in both animal and human studies [9],[10]. However, there are limited TUS systems, and further development of these systems is required to explore the complexities of ultrasound-induced neuronal activities.

1.2 Preliminary Work by Biomedical Ultrasound Group

To target the deep brain accurately and precisely through an intact skull for transcranial ultrasound stimulation, a hemispherical phased multi-element array was developed in the Biomedical Ultrasound Group. A phased multi-element array can be driven to electronically correct variations induced by the skull at the individual element level. The array contains 256 individually controlled elements, randomly arranged to suppress side lobe formation and to optimize steering and the focal spot shape. A Figure of the TUS helmet can be seen in Figure 3. The multi-element TUS system developed by BUG is an MR-compatible helmet that is secured on top of the subject head using custom-made 3D-printed face masks. The subject will receive a pre-therapeutic MRI or CT scan for therapeutic planning purposes. If the location and acoustic properties of the skull are known, BUG has simulation software that will calculate the phase shifts required to correct for skull variations.

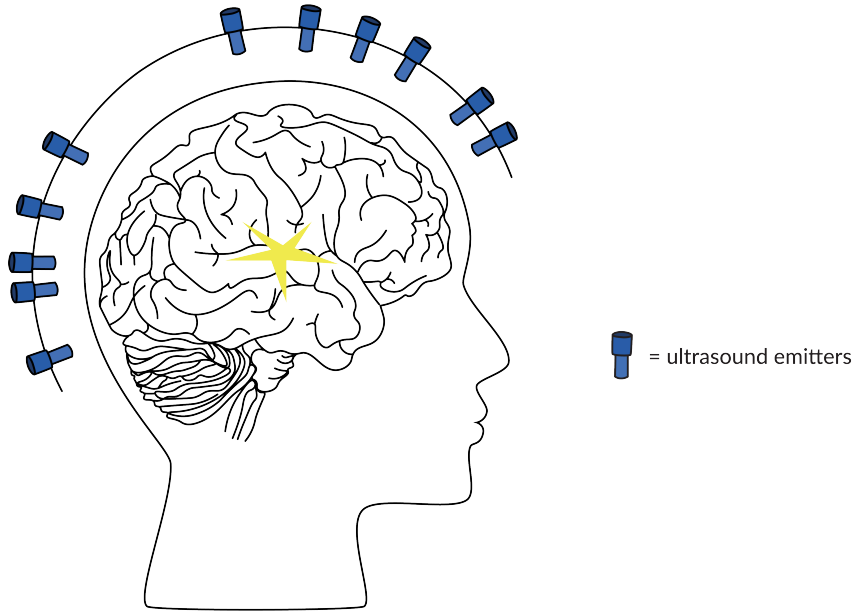


Figure 3: Sagittal plane slice of Biomedical Ultrasound Group TUS helmet.

1.3 Problem Statement

The TUS helmet relies on sending multiple low-intensity ultrasound pulses through the skull with specific time delays such that they constructively interfere at chosen foci in the brain. However, bone is a highly reflective and absorptive material that limits the transmission of acoustic waves. Additionally, the skull is a heterogeneous material, varying in thickness and acoustic properties, so different locations around the skull have different effects on the acoustic waves. Passing acoustic waves through the skull results in a distorted acoustic field and thus distorted foci. A figure of skull-induced distortions can be seen in Figure 4. The aberrations can be directly accounted for if the acoustic properties of the skull and the location of the skull relative to the array are known. Changes to the field must be measured or modeled accurately enough to correct the time of flight delays to maintain the original target. Calculating the time of flight delays can be done through various approaches, including both invasive and non-invasive methods.

A non-invasive method developed for the TUS helmet in BUG uses a pre-therapeutic

CT/MR image of the skull to calculate the skull induced phase-shifts. However, the TUS helmet must be in the coordinate system of the subject’s pre-therapeutic CT/MR image. The custom-made 3D-printed face masks are fastened securely to the device to control the subject’s location in a known way. However, the subject can still be several mm off the pre-operative planned position, which is enough to change the time of flight delays needed to maintain the same target foci. The goal of this project was to create, test, and validate an algorithm to acquire data from the TUS helmet, form a point cloud of the head outer surface, and register the point cloud to the pre-therapeutic image. The solution workflow is described in Figure 5. The algorithm must output a transformation matrix that transforms the TUS helmet into the alternate modality image coordinate system.

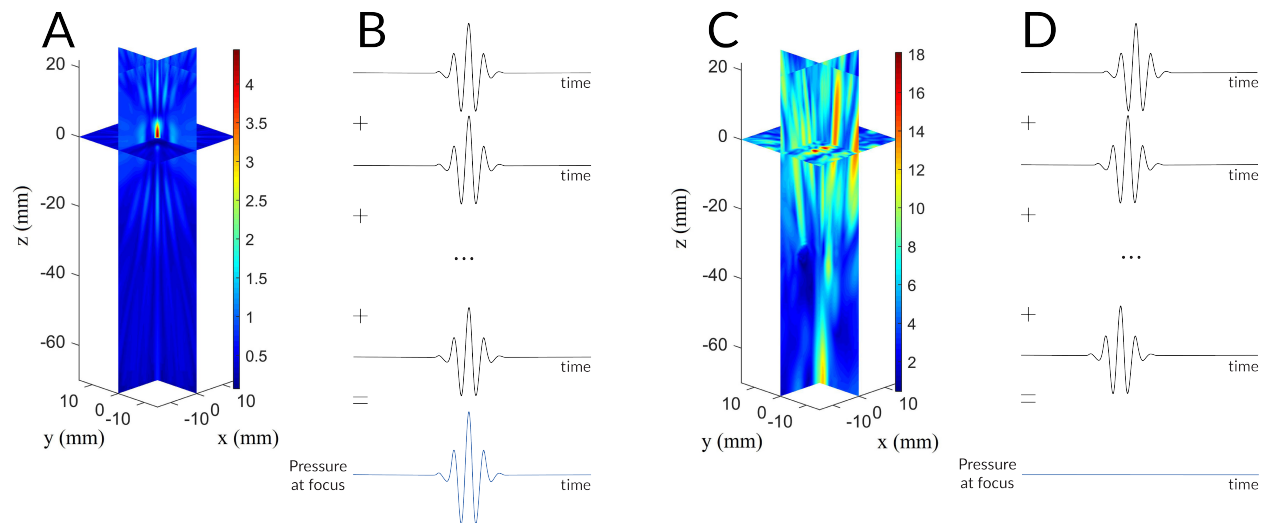


Figure 4: (A) Acoustic field from a phased-array probe without the skull present, showing conventional focusing through constructive interference (illustrated in B). (C) Acoustic field from same array with skull present, showing that the pulses get shifted in time and no longer add up constructively at the geometric focus (illustrated in D) reproduced from a full wave-simulation evaluation [11].

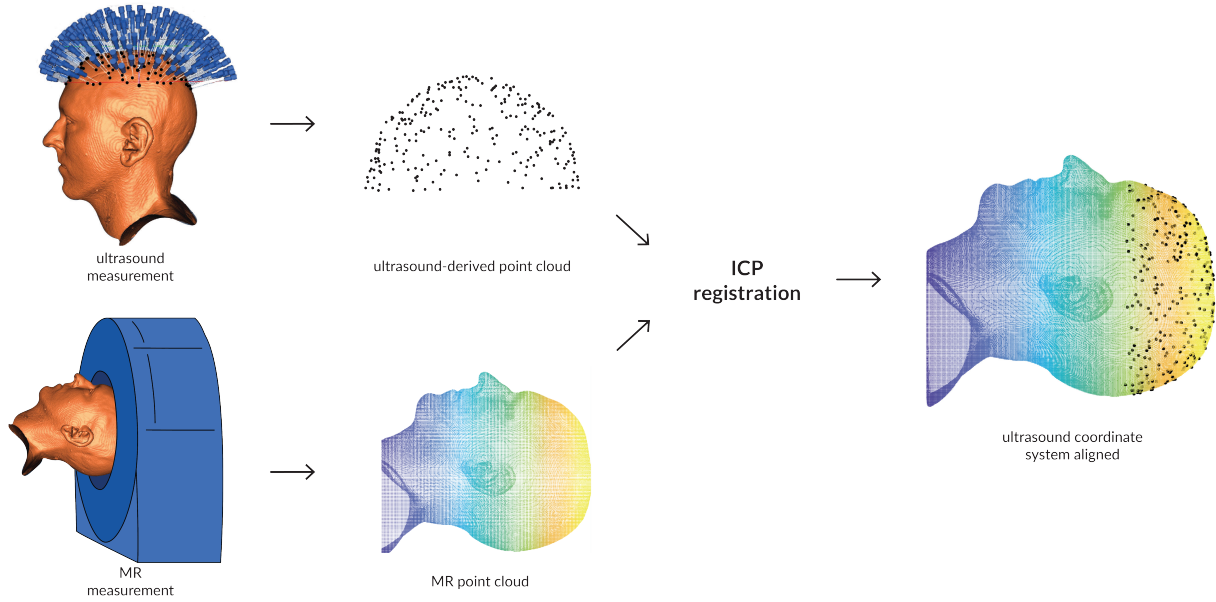


Figure 5: Workflow describing proposed solution to determine transducer helmet location with respect to an MR image. The point clouds are first extracted from the medical imaging data, ultrasound and MR or CT, then they are registered using an ICP registration method.

1.4 Related Works

Iterative closest point (ICP) registration is the current gold standard for rigid body registration of 3D shapes first described by Besl & McKay in 1992 [12]. It works through finding the closest point pairs between two geometric 3D bodies and moving to minimize the mean-square distance metric. Due to the simplicity of the registration, the algorithm may converge on a local minimum. So for accurate registrations a good initial starting guess and shapes with distinct features and a lack of symmetry are recommended.

Fieten et al. in 2009 developed an adaptation of the ICP algorithm that uses a random-weighted ICP, R-ICP, method to register ultrasound-derived points on the skull with the bone surface extracted from a CT image for robotic skull milling [13]. The approach first pre-aligned the points and CT image using a rough guess followed by a random-weighted surface-based ICP registration. The method has the advantage of avoiding getting trapped in a local minima because of the random weighting factors, which are included after the point correspondences and prior to the rotation and translation parameter optimization as

seen in Section 2.3. Once the R-ICP algorithm has terminated, Fieten et al. use this solution as the starting point for a classic ICP method.

Comparisons between the R-ICP method and the classic ICP registration showed that the difference between a ground truth and the registration were 1.30 ± 0.49 mm and 5.66 ± 2.75 mm for R-ICP and classic ICP registrations respectively. The ground truth used was a visual fiducial marker-based registration. The difference between each method is partially due to brute force, since this method first uses an initial starting guess, runs the R-ICP algorithm 100 times and then uses the best solution as the initial starting point for the classic ICP. The method practically runs the ICP algorithm over 100 times with different initial guesses to find the best solution which would help to find the global minimum error. To minimize computation time, Fieten et al. pre-computed a hybrid kd-tree to create a subset of points accessible for point correspondence. Integrating the visual feedback, a few points on the surface of the skull are determined and registered with the CT-derived points. Error in the registration method could come from the points determined from the skull point acquisition method. The Fieten et al. method would be both time and computationally expensive and the visual tracking system for determining ultrasound points has a considerable impact on their registration accuracy.

In 2016, O'Reilly et al. used a hemispherical array to rigidly register the CT-derived skull outer surface to the ultrasound-derived bounding surface for correction of skull-induced phase variations during transcranial focused ultrasound [4]. The hemispherical array was driven using a single-element single cycle burst and the same element receiving, averaged over five acquisitions. They mathematically created bounding spheres centered around the element with radius equal to half the time of flight over the speed of sound. This resulted in 96 bounding spheres that they aligned the CT-derived skull outer surface to.

First, the algorithm solves for the CT-derived data translation to the center of the helmet by minimizing the difference between the distance between closest skull points and elements

and the pulse-echo distances. Then it loops through perturbations of constrained movements within ± 3 mm and $\pm 3^\circ$ in every direction until it minimizes a cost function that contains a penalty if the skull lies closer to the elements than the bounding surface. The resulting skull location was within 1 mm and 1° of the ground truth. The ground truth was attained from multiple CT images where the location of the skull and ultrasound helmet can be seen and the absolute displacement and rotational errors were calculated. The O'Reilly et al. method is a simple algorithm for finding the points on the skull and did not use triangulation of signals received from other elements, so some error may arise from using so few bounding surfaces. Additionally, the initial guess must be within 3° in any axis off of the true position due to the constrained nature of the rotational solver.

In 2018, Crake et al. used the same method for 3D skull localization laid out in the O'Reilly & Jones paper for a hemispherical array, but with a 48 element array showing errors of 2 mm and 2° extracted from the ground truth extracted directly from an MR image (± 0.5 mm/ 0.5°) [14]. Then in 2021, Deng et al. also used the O'Reilly et al. method, but instead with an constrained solver of $\pm 10^\circ$ and ± 3 mm along each axis, resulting in registration errors of 1.4 ± 0.4 mm and $2.1 \pm 0.2^\circ$ derived from a optical landmark-based ground truth method [15].

In 2022, Lu et al. used pulse-echo measurements from a hemispherical array to derive points on the skull and then used a classic ICP algorithm using this ultrasound-derived point cloud and a CT-derived skull surface [16]. They didn't compare to a ground truth image and instead used the ICP output RMSE as their metric of comparison. Using this metric, they reported an error of less than 2.5 mm for all registrations. However, this metric may not be an accurate representation of the error, as the rotational error could be much larger since it has a minimal effect on the distance between elements and skull surface. Additionally, they used a classic ICP registration method without multi-start, which depending on the initial start guess may have converged to a local minimum.

1.5 Objectives

The objective of this project was to design an algorithm that would perform an ultrasound-based skull registration for the purpose of skull localization. The process included:

- Designing and creating data acquisition scripts, utilizing each element of the TUS hemi-spherical array
- Processing the pulse-echo data from the hemi-spherical array to recover useful information
- Deriving the mathematics to reconstruct a point cloud of the object in the ultrasound field's outer surface
- Creating a function to register the ultrasound-derived point cloud to an alternative 3D image of the object from either a CAD file or for human experiments, a pre-therapeutic CT or MRI image
- Testing scripts experimentally to determine the boundaries and limitations of the algorithm
- Validating scripts with a ground truth MR image
- Verifying that the matrix can be used to transform the ultrasound therapy plan to derive the necessary phase and amplitude corrections to correct for skull-induced aberrations to the acoustic field

CHAPTER 2

REGISTRATION ALGORITHM DEVELOPMENT

This chapter describes the mathematics used to reconstruct the ultrasound-derived point cloud, the development of acquisition scripts, and the registration method created for the registration algorithm.

2.1 Ultrasound Point Cloud Reconstruction

To derive the points on the surface of the skull given ultrasound data from the TUS helmet, some assumptions needed to be made. The first being that the path the acoustic wave travels forms an isosceles triangle with the transmit and receive elements and the reflection point on the skull, shown in Figure 6. This assumes spherical spreading wavefronts from the elements. The sides of the isosceles triangle are calculated from the time of flight data and speed of sound of the water, which is calculated for each ultrasound measurement. The isosceles triangle assumption degrades with increasing distance between transmit and receive elements, so only elements within 3 cm of the transmitting element were used to map the outer surface. The third side of the isosceles triangle is also known since the location of the transmit and receive element are known, making it straightforward to calculate the height of the isosceles triangle.

Knowing the halfway point between the two elements and the height of the triangle results in an infinite number of points as the triangle can be rotated 360 degrees. A second assumption was introduced: the triangle is pointing in the direction normal to the surface of the helmet. This is a reasonable assumption because the helmet was designed to replicate the average skull but enlarged. The helmet was modelled using a tri-axial ellipsoid function with defined radii, illustrated in Figure 7. The normal vector to a surface is defined by the grad of the parameterized function of that surface. The inward-pointing unit normal vector

is simply the negative of that vector divided by the magnitude of the vector. The x,y,z points of the interpolated element were substituted into the vector and the height of the isosceles triangle was multiplied. Finally, the vector was added to the position of the interpolated element to find the x,y,z coordinates of the skull outer surface.




Figure 6: The Figure shows the assumed pathway the acoustic wave travels from the transmit element (Tx Element i) and the receive element (Rx Element j) forming an isosceles triangle where $d_i = d_k$. This assumption is the base for determining the height of the triangle that is formed between the two elements and the reflection point on the surface of the skull.

2.2 Acquisition Scripts

To acquire the data from the TUS helmet, a script that performs a synthetic transmit aperture was written. The script performs a single element transmission and all elements receiving image sequence, transmitting from each of the 256 elements in turn. Each element is driven individually with a 1 cycle tri-state pulse at 20 V with a transmit frequency of 0.554 MHz. The acquisition time is set to 250 μs with a pulse repetition interval of 500 μs . Then, each element is transmitted 32 times, although this number can be changed by the

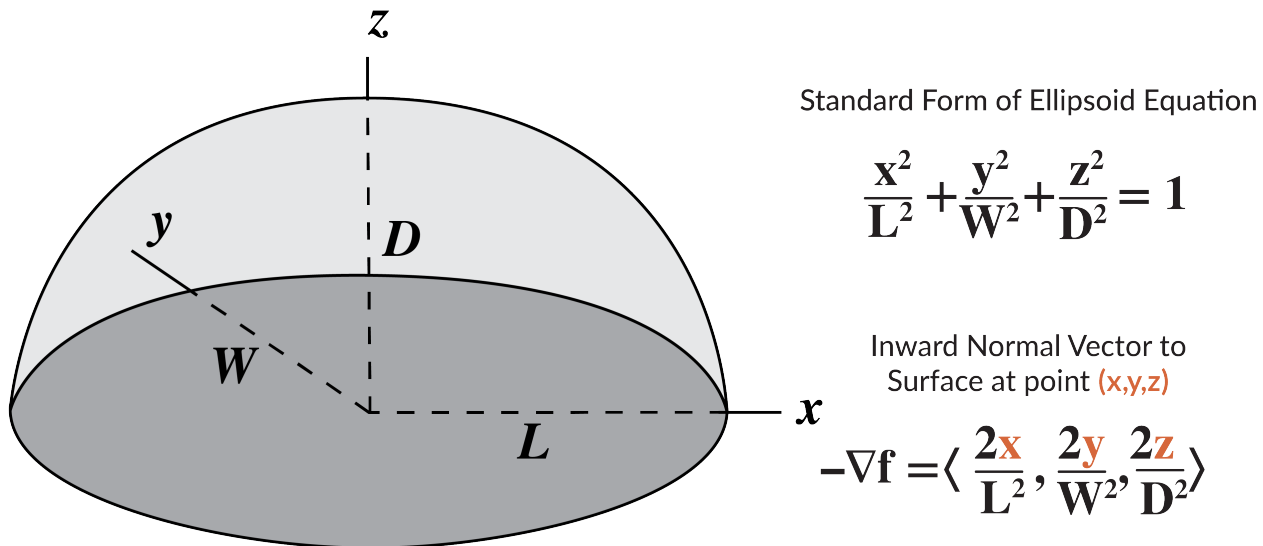


Figure 7: The inward normal vector can be worked out to be the negative of the parameterization of the standard form of the ellipsoid equation, since the radii of the tri-axial ellipsoid, D , W , and L from the figure, are known.

user, and the result is averaged for the final dataset. The data from transducers outside a 3cm radius of the transmitting element is not used in the reconstruction of the point cloud to uphold the isosceles triangle assumption.

Given that there are many different documented ways for finding the time of arrival of a signal in the literature, a set of experiments was performed. Two elements were positioned in a tank at a measured location pointing at each other as seen in Figure 8. Then, a pulse with the same parameters as the imaging parameters was sent on one of the elements and the second element recorded the signal. Four sets of transmit-receive element pairs and multiple distances, between 4 - 8 cm, were tested. This distance was chosen because it is the expected round trip distance that will be seen in skull experiments. The time series waveforms were converted to distance by multiplying by the speed of sound of the water which was tested after each experiment.

Three functions were created in MATLAB to pick the time of arrival, the yellow star in Figure 9. All signals were clipped to remove the initial EM interference. The first method,

the cumulative sum method, picked the time of arrival as the point when the cumulative sum of the signal reached a specific threshold, 5% in this case. The second method, the derivative method, picked the time of arrival as the largest change in the differential of the smoothed, enveloped signal. The third method, the minimum AIC method, picked the time of arrival using the minimum value of the Akaike Information Criterion function which uses the log likelihood of a sliding window to calculate a value that is representative of the amount of variation between signal segments within the window. The results of the experiments were plotted in Figure 10. The minimum AIC method predicted the time of arrival with closest alignment to the calculated distances in the laboratory experiments. Therefore, this function was the method used in the point cloud reconstruction.

Using the acquisition scripts, data processing, and derived reconstruction mathematics, an ultrasound-derived point cloud from the surface of an object was able to be mapped. The ultrasound-derived point cloud from a 3D-printed PVA skull phantom, described in Section 4.1, is visualized in Figure 11.




Figure 8: The transmit-receive element pair are setup such that they point directly at each other and are held at measured distance between 4-8cm apart.

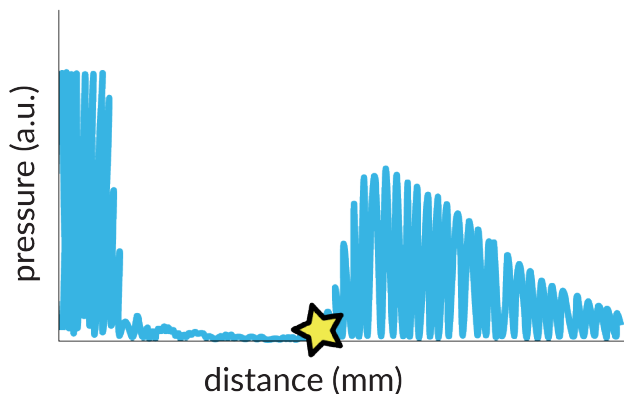


Figure 9: An example of a recorded ultrasound signal is plotted on a graph with pressure on the y-axis and distance on the x-axis. The time of arrival labeled with a star is the data point that the time of flight scripts were designed to find.




Figure 10: The calculated vs. measured distance for each method for picking the time of arrival showing the minimum AIC in greatest alignment with the measured distances.

2.3 Registration Method

To register the ultrasound-derived point cloud to the pre-therapeutic MR/CT image, an iterative closest point (ICP) registration method was employed using a built-in MATLAB function, `pcregistericp` [17]. The function requires one moving point cloud and one fixed point cloud and outputs a transformation matrix that moves the moving point cloud to the location of a local greatest alignment with the fixed point cloud. An example of the fixed and moving point clouds can be seen in Figure 12. The ICP registration is one of the oldest and most classic methods for rigid body registration first described in 1992 by Besl & McKay



Figure 11: The ultrasound-derived point cloud from the surface of a 3D-printed PVA skull model showing the point cloud reconstruction method qualitatively looks successful.

[12]. It works by alternating between closest point query (determining closest point pairs) and determining a rotation and translation that moves to minimize the distance between the point pairs. The ICP method converges when the root mean squared error (RMSE) of the Euclidean distance between aligned points is within a set tolerance value provided or if the maximum number of iterations set by the user is reached. A figure of the method is described in Figure 13 and the mathematics are more rigorously described below.


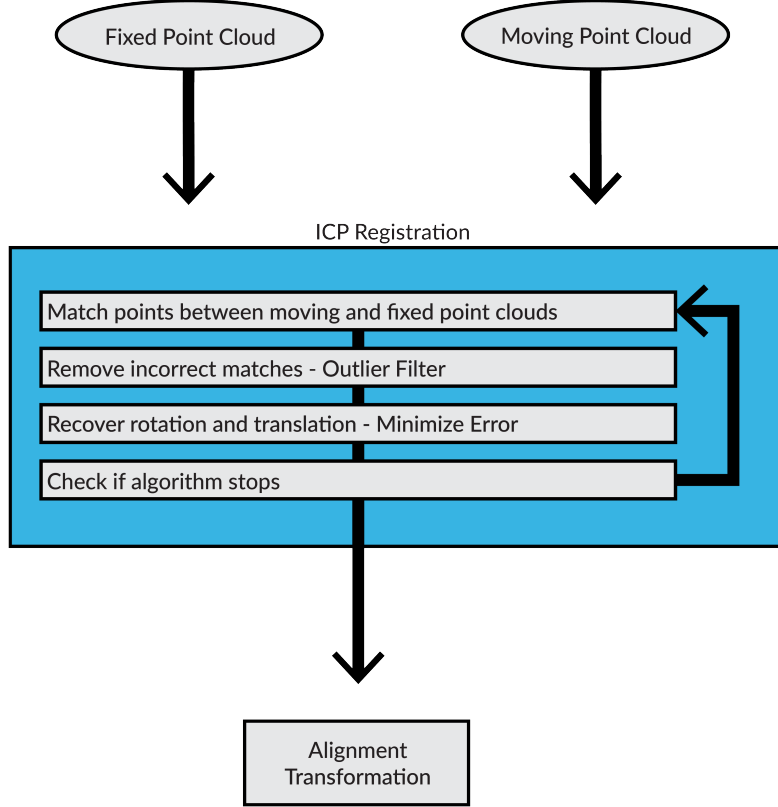


Figure 12: An example of the point clouds passed to the ICP function. The fixed point cloud was derived from the mesh of the CAD file used to print the skull model. The ring around the base is the extruded polyline used to close the bottom surface.



The MathWorks, Inc. (2014-2021)

Figure 13: The ICP registration workflow taken from the MATLAB built-in function help page describes how the two point clouds are passed to the function and the function alternates between closest point query and moving to minimize the error [17].

The ICP algorithm works to find a rigid transformation with parameters $\mathbf{T} = (R, t)$ where R and t are the rotation and translation respectively [18]. With fixed and moving point clouds $\mathbf{F} = \{\mathbf{f}_1, \dots, \mathbf{f}_{N_f}\}$ and $\mathbf{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_{N_m}\}$, the error function is defined by:

$$E(\mathbf{T}, \mathbf{F}, \mathbf{M}) = \sum_{i=1}^{N_f} \|(R\mathbf{f}_i + t) - \mathbf{m}_j\|_2 \quad (2.1)$$

where $(\mathbf{f}_i, \mathbf{m}_j)$ are the previously defined closest point pair corresponding points. The ICP algorithm then fixes the corresponding point pairs and finds the values for R and t that minimizes the error. It does this through single value decomposition of the cross-covariance

matrix, \mathbf{C} , between the N_f correspondences:

$$USV^T = \mathbf{C} \quad (2.2)$$

Such that, the rotation matrix can be calculated as:

$$R = VU^T \quad (2.3)$$

Then from the rotation matrix, the translation vector, t , is estimated as:

$$t = \bar{\mathbf{m}} - R\bar{\mathbf{f}} \quad (2.4)$$

Once these are calculated, the point clouds are moved by \mathbf{T} and a new point correspondence matrix is calculated. The error is calculated from equation (2.1) and if the error between two successive movements is lower than a threshold, the algorithm stops. This process is described in greater detail in 3D shape registration [18].

Initially, the ICP method was numerically tested using a point cloud from a perfect ellipsoid. The moving point cloud was defined as the point cloud from the perfect ellipsoid transformed by a known transformation and the fixed point cloud was the original point cloud centered at the origin. The ICP function did not converge on the correct solution. After plotting a few results, it became apparent that the error resulted from the model being too perfectly smooth, so there were not enough features for the algorithm to register the point clouds accurately. The ICP registration requires that the objects to be registered have a certain level of complexity [12]. So, the RMSE of the two point clouds was falling below the tolerance at an incorrect solution. To check that this method works with a head-shaped object, the ICP method was tested using a trimmed skin-mesh derived point cloud from an MR image of a human head and the same point cloud transformed by a known transformation. The skin-mesh from the MR image contains the point cloud of the outer

surface of the person from the shoulders to the top of the head, so the point cloud had to be trimmed to contain only the top of the head replicating the surface that the ultrasound-derived point cloud maps. The resulting solution from this test registration matched the known transformation matrix with a tolerance of 0.2%, suggesting that the human head contains enough structural features for the ICP method to register. An exemplary solution can be seen in Figure 14.

The ICP method created implements a pre-registration step using a rough initial guess prior to the classic ICP registration to avoid getting trapped in local minima. Additionally, a multi-start approach was taken such that the moving point cloud was rotated by 4 initial rotations and then the ICP solution with the lowest error was returned to the user.

This ICP method works really well if the point clouds are identical, but in this setup this is not the case. The ultrasound derived point cloud has 2052 points and the mesh derived point cloud has between 60,000 and 80,000 points. There are errors induced from electronic delays in transmitting and receiving the ultrasound signal which will result in some level of scaling of the ultrasound-derived point cloud. Also, there will likely be noise in the ultrasound-derived point cloud. Additionally, the ICP registration has limitations especially surrounding the initial starting point, since the ICP algorithm works to minimize the RMSE, it may converge to local minima which may not be the global minima. For this reason, these parameters were investigated in a set of numerical experiments, described in the following chapter.

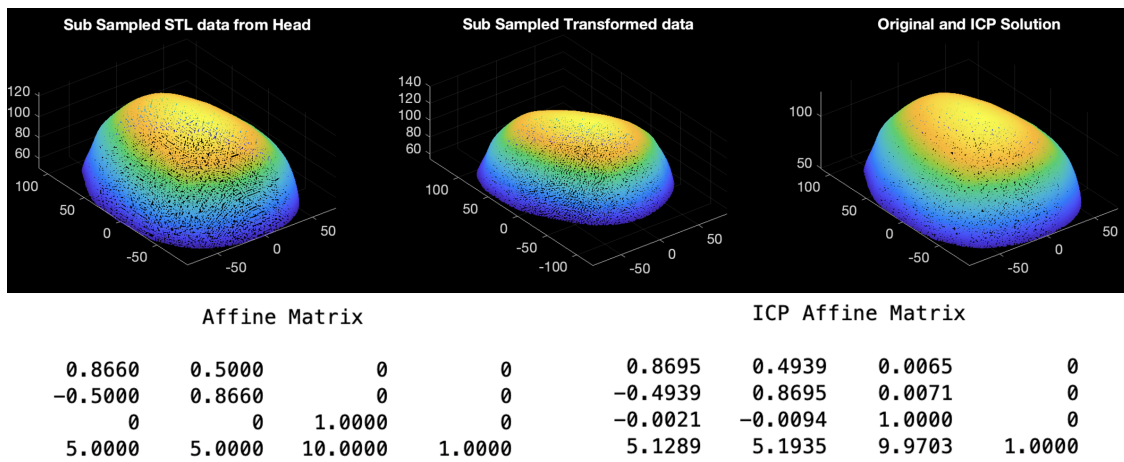


Figure 14: The leftmost plot: skin-mesh derived point cloud from an MR image of a volunteer in it's original coordinate system. The middle plot: the same point cloud transformed by the transformation matrix below the plots on the left. The original data and the transformed data were registered together using the transformation matrix on the right. The rightmost plot: the original point cloud plotted with the image transformed using solution from the ICP function.

CHAPTER 3

NUMERICAL EXPERIMENTS

This chapter describes the numerical experiments performed to test the registration algorithm and summarizes and discusses its limitations. The numerical experiments test two overarching difficulties seen in a real world setting: the initial starting point for the moving point cloud not being perfect and the two point clouds not being identical. To test the initial starting point, the moving point cloud's initial starting point was rotated and translated further away from the fixed point cloud. To test the possible differences between the point clouds, noisy, down-sampled, and scaled data was simulated independently.

3.1 Methods

Considering the limitations of the ICP algorithm, it was robustly tested to quantify its limits when used in the context of the TUS experiments. Three point clouds derived from three MR images of different volunteers were used along with different parameters following the workflow described in the Figure 15. The fixed point cloud was the original skin-mesh derived point cloud from the MR image and the moving point cloud was this point cloud after changing a parameter and transforming it by a random transformation which can be seen in Figure 16. The random transformation was made up of a random rotation of up to $\pm 10^\circ$ in every axis, sampled with a uniform distribution, and a random translation with up to ± 1 cm in every direction. These limits were chosen based on an estimation of the maximum position error likely to be seen in a real world experiment.

For each variation of every parameter, the moving point cloud was transformed 50 times by a different random transformation and passed through the ICP function. This was repeated for each of the three datasets. From the output transformation from the ICP function, the Euler angles of the rotation were calculated [19]. The root mean square error between

the angles of rotation and the root mean square error in translation was calculated as a metric to determine the "goodness" of alignment. In order to correct for the variations to the acoustic field induced by the skull, it is required that we know the skull's location within about 1 mm of its translational location and 3° of its rotational location. Each of the results were compared with these chosen thresholds.

A function to test the ICP function was written. The user inputs the point cloud to be tested, rotational angle and rotational point for which to trim the mesh-derived point cloud, uniform scaling tested, power of the noise, down-sampling step size, and optionally a transformation matrix if a random transform is not desired.

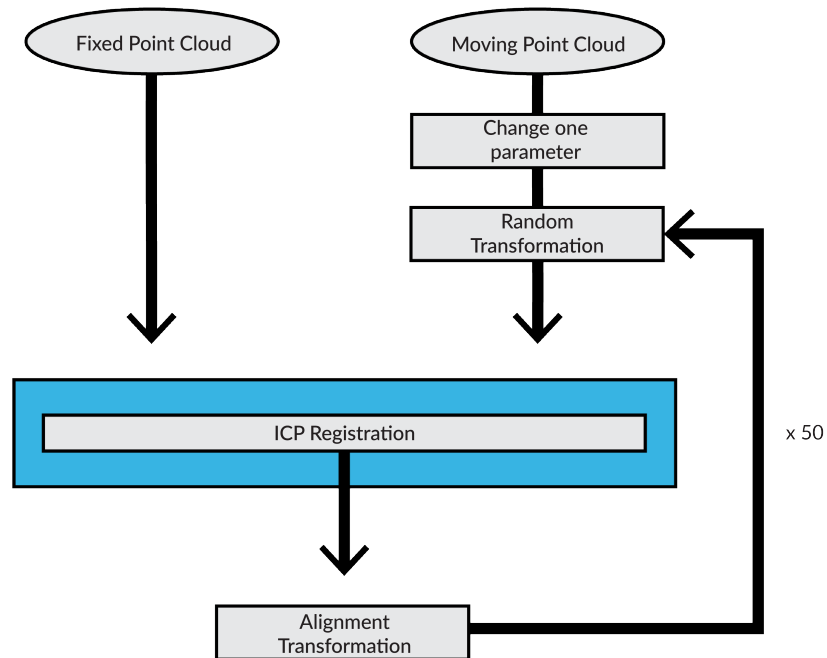


Figure 15: The numerical experiment workflow describing the process of testing each parameter variation.




Figure 16: An example of the point clouds passed to the ICP function for the numerical experiment. The moving point cloud in this example is down-sampled and noisy.

3.2 Number of Point Pairs

In a real world experiment, the ultrasound-derived point cloud will be made up of 2052 points due to the limited number of ultrasound transducers, so the number of points for the moving point cloud was varied to see how the ICP algorithm responded. The function as described above was used for each of the three datasets for each down-sampling variation. The down-sampling variation was passed to the function 50 times for each of the datasets. The down-sampling resulted in the moving point cloud containing between 500 to 4000 points. This range was chosen to simulate the number of points that the ultrasound-derived point cloud contains. The other parameters passed to the function were a uniform scaling of one and zero noise. Each time the moving point cloud is down-sampled, a random subset of the points are picked without replacement sampled uniformly at random.

The RMSE between Euler angles from the applied and ICP transform remained below 3° , which was deemed an acceptable error, as can be seen in Figure 17. However, the translation RMSE was considerably larger (Figure 18). The median for most of the range was above 4 mm. The ICP RMSE was relatively constant across the number of point pair range (Figure 19). This metric value isn't helpful for direct comparison. For the remaining experiments, the moving point cloud was down-sampled to approximately 2000 points.

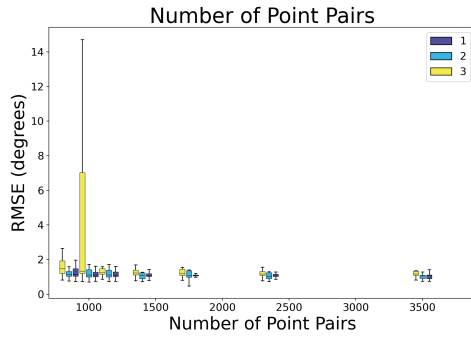


Figure 17: The RMSE between the angles of rotation between actual and ICP transformations vs. the number of points that the moving point cloud contains showing minimal change in the range of interest.




Figure 18: The RMSE of the translations between actual and ICP transformations vs the number of points that the moving point cloud contains showing seemingly high translational errors.

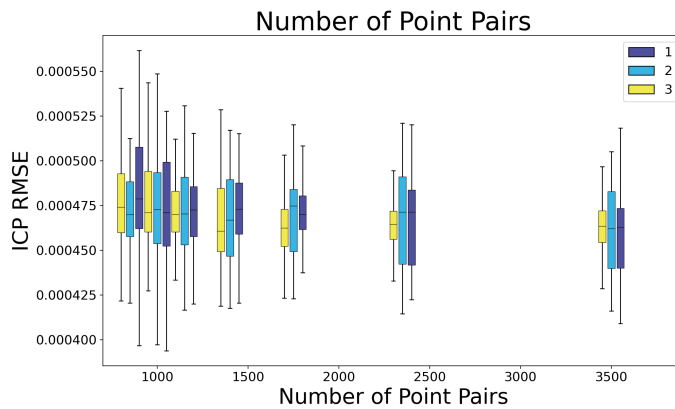


Figure 19: The ICP RMSE vs the number of points that the moving point cloud contains showing a consistent trend across all point pairs of interest.

3.3 Noise

In a real world experiment, it is likely that there will be noise added to the data from various sources, so random Gaussian noise with varying power was tested. This was done using a MATLAB built-in function, `randn`, to generate a matrix with the same size as the moving point cloud with random values. The function works by generating data that follows

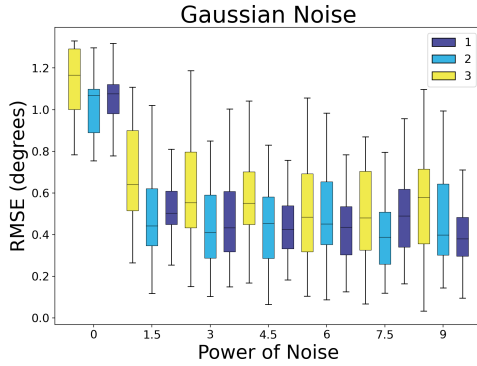


Figure 20: The RMSE between the angles of rotation between actual and ICP transformations vs the power of Gaussian noise showing that there is not a significant change across noise powers, except that there is a slight decrease in error with added noise.




Figure 21: The RMSE of the translations between actual and ICP transformations vs the power of Gaussian noise also showing that there is a slight decrease in error with added noise.

the standard normal distribution with a mean of zero and a standard deviation of one [17]. The noise matrix is multiplied by the square root of the power of the noise passed as a parameter to the testing function. Then, the noise matrix was added to the original moving point cloud. Noise power levels between zero and nine were used.

The RMSE between Euler angles from the applied and ICP transform remained below 3° and remained constant across all noise power levels (Figure 20). However, again the translation RMSE was larger (Figure 21). The median for every noise level was above 1 mm. The ICP RMSE increased with increasing noise (Figure 22). This is a metric that measures the relative error between the points and since the moving point cloud became noisier, this is not representative of the true positional error.

3.4 Uniform Scaling

The ICP testing function was used to iterate through variations in uniform scaling, since real world electrical delays in the transmit and receive signal processing may result in some

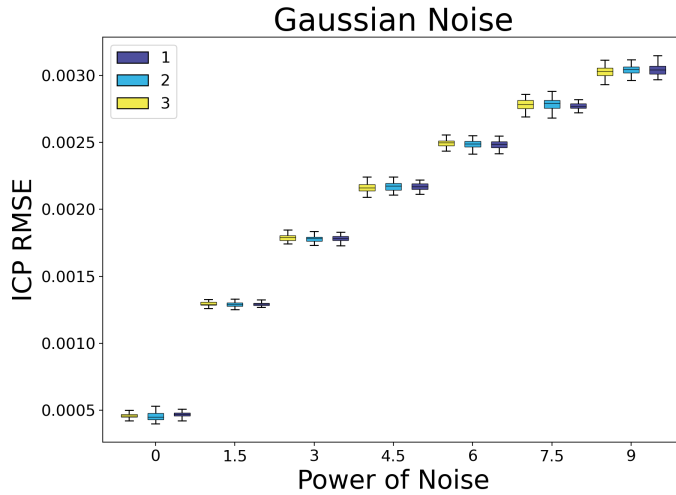


Figure 22: The ICP RMSE vs the power of Gaussian noise showing an increase in error with increasing noise, which is as expected since the noise will increase the average distance between the point pairs.

scaling issues. The moving point cloud was scaled in a range from $1/2$ to $3/2$ of its original size. Because the moving point cloud was neither centered at the origin nor a symmetric shape, dilation/contraction was not simply a matter of multiplying by a scaling factor. Scaling involved translating the center of the object to the origin. Then, applying the scaling factor and finally, translating back to its original point. Notably, scaling is not considered a rigid body transformation, since the shape may not have the same shape as after the transformation if the shape is not symmetric. In this case, the skull is definitely not symmetrical.

Although this set of experiments is not reflective of truly what happens when electronic delays are introduced to the signal, a limitation of the ICP algorithm was revealed through this exploration. The ICP algorithm produced large errors in every metric of comparison (Figures 24 - 26). After plotting an example of the resulting alignment, which can be seen in Figure 23, it became clear that the ICP algorithm was converging on these incorrect solutions because after every iteration the algorithm picks new closest point pairs without any guidance on which point pairs wouldn't be feasible pairs physically. One possible solution

to this would be to use 1-1 point correspondence derived from experimental US data and a simulation of the array using a reasonable assumption of its location or to limit what subset of the model points can be used for closest point pairs which would also significantly improve the speed of the ICP algorithm. Additionally, the electronic delays that likely add some level of scaling should be removed prior to registration.

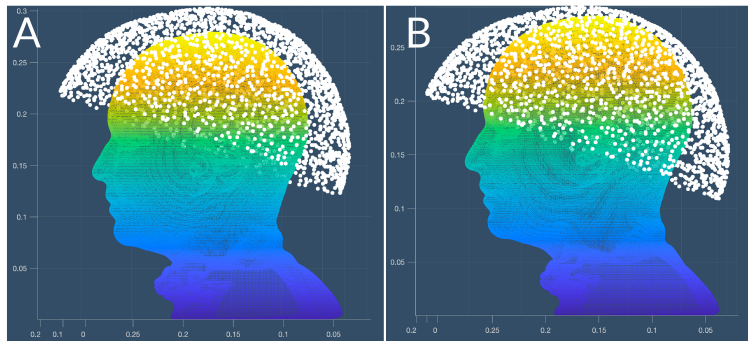


Figure 23: (A) scaled point cloud and (B) scaled point cloud moved to ICP algorithm solution showing the convergence to the incorrect solution.

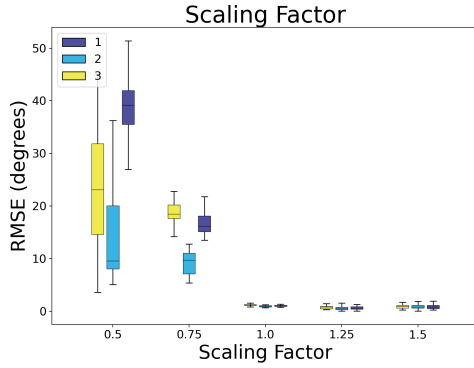


Figure 24: The RMSE between the angles of rotation between actual and ICP transformations vs the scaling factor shows the large error when the moving point cloud is scaled.

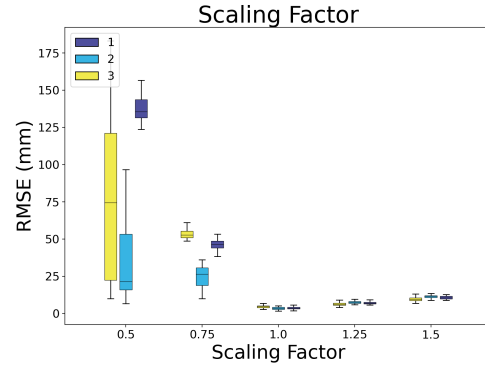


Figure 25: The RMSE of the translations between actual and ICP transformations vs the scaling factor, also showing error as the moving point cloud is scaled.




Figure 26: The ICP RMSE vs the scaling factor also shows error as the moving point cloud is scaled.

3.5 Rotation

In the real world experiments it is possible that the subject will be placed into the TUS helmet with various starting positions, so the next two sections are the results of testing both rotations and translations independently.

The starting position of the moving point cloud was rotated by between $\pm 30^\circ$ for each

axis without any translation. However, for this parameter the ICP registration was only tested one time per parameter variation per dataset. Figure 27 illustrates the three axes of rotation and how they relate to the head. The scaling factor was set to one, noise to zero, and down-sampling such that the moving point cloud had 2000 points. Then, the difference between the Euler angle for the respective axis of translation was calculated. This is used as the metric for comparison in Figure 28. The data in Figure 28 shows that as the image is rotated further from the true location, the standard deviation of the error gets larger. The error is quite large compared to the other rotational errors seen from other parameters, ranging between 0 to 12° of error for as little as a 10° rotation in starting position.




Figure 27: Figure showing the three axes of rotation with respect to the subject's head.

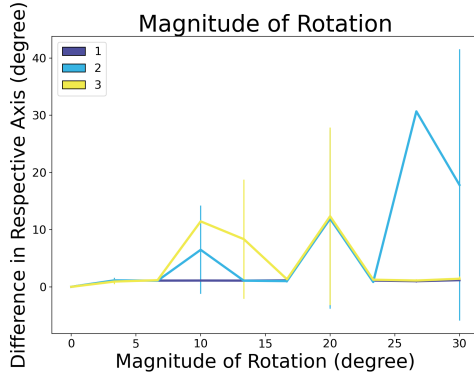


Figure 28: The difference between the angles of rotation between actual and ICP transformations vs degrees of rotation. The lines represent the standard deviation of the difference. As the initial starting point's rotation increases, the standard deviation of the rotational error increases.




Figure 29: The RMSE of the translations between actual and ICP transformations vs degrees of rotation shows that with increases rotation the translational error increases.




Figure 30: The ICP RMSE vs degrees of rotation shows the ICP error increases with increasing rotation.

3.6 Translation

Finally, the moving point cloud was translated between ± 1 cm in a random direction. The magnitude of translation was generated randomly and passed through the ICP algorithm

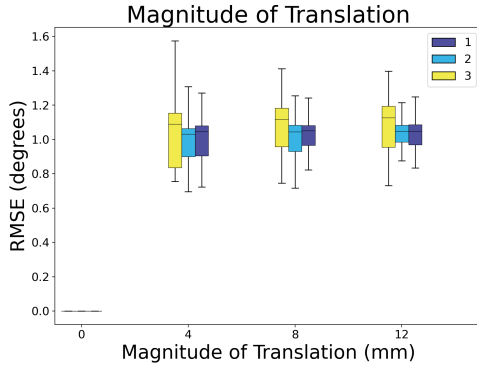


Figure 31: The RMSE of the rotations between actual and ICP transformations vs translation of moving point cloud showing the rotational error not changing significantly with increasing translation in starting position.




Figure 32: The RMSE of the translations between actual and ICP transformations vs translation of moving point cloud showing the translational error also not changing significantly with increasing translation in starting position.

50 times per dataset. The rotational RMSE in Figure 31 is within the 3° threshold, but the translational RMSE in Figure 32 was not within the 1 mm requirement.

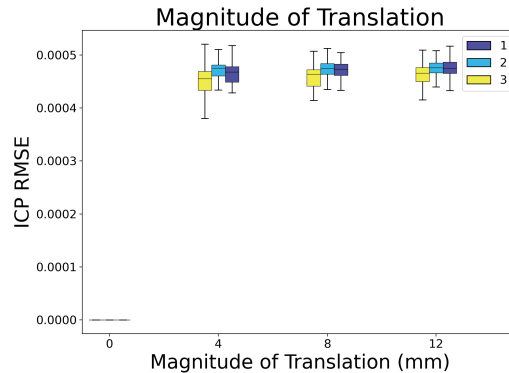


Figure 33: The ICP RMSE vs translation of moving point cloud showing minimal change in ICP error with increasing translation in starting position.

3.7 Summary of Results

After exploration of the parameters, the data showed that down-sampling, noise, and translating the initial guess had no significant effect on the performance of the ICP algorithm at the range of interest for this use case. With the exception that with added noise,

the rotational and translational error dropped slightly. Conversely, the ICP algorithm converged on an incorrect solution when the moving and fixed point clouds were not the same scale. There are a couple discussed solutions to this problem: forced 1-1 point correspondence, previously defined k-d trees, or removing scaling altogether from the data. Also, the data suggests the ICP is not as robust to rotational movements as it is to translational movements. This is due to the approximate rotational symmetry of the shape of the skull. Although within the given range of possible subject movement that is expected, it is likely that rotations will stay below 10° in every axis and the rotational error for this range still remained lower than the chosen rotational threshold.

For the phase shifts induced by the skull to be corrected, the skull location needs to be known within a rotational limit of 3° and translational limit of 1 mm. Unfortunately, through all these numerical experiments, the results showed the translational error was above these values suggesting that the current version of the ICP algorithm doesn't satisfy the requirements for image registration. Although, with the process for testing streamlined, new versions can be tested quickly to determine the appropriate version moving forward.

Interestingly, the returned ICP squared distance metric (RMSE) was not a good metric to use to determine the "goodness" of alignment, it was sensitive to noise, partial overlaps, and not sensitive to rotational error. Using the ICP error metric may lead to making incorrect assumptions about registration accuracy. Overall, the ICP algorithm is highly sensitive to the initial starting point, the shape and symmetry of the object being registered, and scaling differences between objects. Researchers that use this method for image registration should consider these limitations and the error that may result from using it.

CHAPTER 4

EXPERIMENTAL VALIDATION USING 3D-PRINTED PVA PHANTOM

This chapter describes the process of validating the registration algorithm in a real world experiment using a 3D-printed model of a head. The 3D-printed model of a head was created, fixed into the TUS helmet, and MR imaged as a ground truth for the ICP algorithm. The ICP algorithm solution for the location of the skull was 2° and 2.9 mm off the ground truth location.

4.1 Methods

A 3D-printed Polyvinyl Alcohol (PVA) model of the skull was used to validate the ICP algorithm using a ground truth. The model outer surface was designed using a skull template derived from MR and CT images of healthy human subjects (mean age 65) and the inner surface was modelled using the standard brain defined by the Montreal Neurological Institute (MNI) [20], [21]. These surfaces can be seen in the Figure 34. In CAD software, the skull inner surface was cut from the skull outer surface. The resulting skull shape was printed using PVA. To compare the location derived from the ICP algorithm with a ground truth, the skull phantom was fixed to a watertight lid within the TUS helmet and the entire unit was taken to the Nuffield Department of Neurosciences at Oxford to get MR Imaged. An image of the kit inside the MR room is shown in Figure 35. Since the 3D-printed model was made from PVA plastic which has a distinct lack of protons to excite, the model and the kit are quite difficult to make out in the localizer MR images. The images were taken with a large object with high MR contrast which made displaying the skull model slightly harder, so the MR images were thresholded for better displaying. A slice of both the original and processed MR image is displayed in Figure 36. Directly following the MR imaging sequence,

an ultrasound imaging measurement was taken.

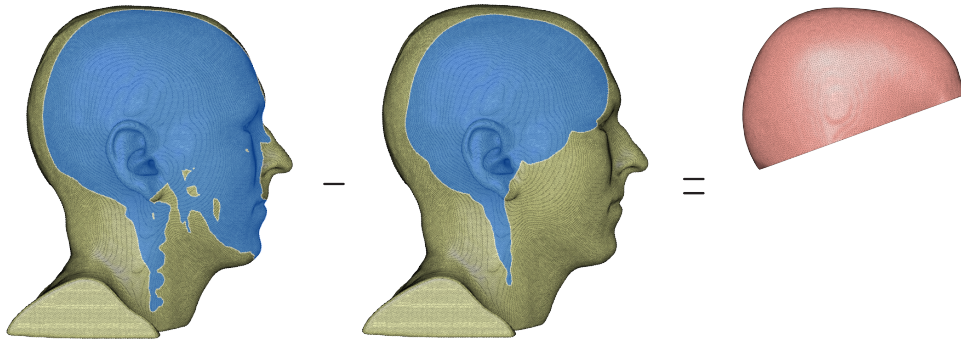


Figure 34: Modelling process for creating the skull phantom from the MNI brain and average skull overlaid on a subject MRI image for reference.



Figure 35: The TUS system in the MR room at Oxford. Unfortunately, we did not get an image of the skull model, so the object seen in the TUS system is half a watermelon.




Figure 36: (A) A slice of the original MR image of the phantom in the TUS helmet, imaged with a rectangular object with high MR contrast. (B) To visualize the phantom and helmet better, a thresholded slice of the MR image of the phantom in the TUS helmet. The phantom is outlined in yellow and the helmet is outlined in red.

4.2 Results

The TUS helmet and skull model CAD files were manually overlaid on the MR image to find the transformations, \mathbf{T}_1 & \mathbf{T}_2 respectively, that moved the files into the MR image space. Figure 38 compares the location of the ultrasound-derived point cloud with the skull model prior to any processing. The ground truth transformation to move the TUS helmet into the skull model coordinate system was found to determine $\mathbf{T}_1 \rightarrow \mathbf{T}_2$:

$$\mathbf{T}_1 \rightarrow \mathbf{T}_2 = Inv[\mathbf{T}_1] * \mathbf{T}_2 \quad (4.1)$$

There was a large discrepancy between the starting position for the ultrasound-derived point cloud and the reference as shown in Figure 38. Consequently, an initial 90° flip about the z-axis was implemented as a rough initial estimate, similar to the process used for the head experiments. The 90° flipped ultrasound-derived point cloud was passed to the multi-start ICP algorithm as the moving point cloud and the skull model as the fixed point cloud. The multi-start ICP algorithm took 10 seconds to converge to a solution and the result can be seen in Figure 39 with the following transformation matrix:

$$R = \begin{bmatrix} 0.0003 & -1.0000 & 0.0063 \\ 0.9990 & 0.0006 & 0.0449 \\ -0.0449 & 0.0063 & 0.9990 \end{bmatrix}, t = \begin{bmatrix} 0.0037 \\ -0.0018 \\ 0.0026 \end{bmatrix} \quad (4.2)$$

The transformation describing $\mathbf{T}_1 \rightarrow \mathbf{T}_2$, the ground truth, is the following:

$$R = \begin{bmatrix} 0.0476 & -0.9988 & 0.0068 \\ 0.9988 & 0.0477 & 0.0070 \\ -0.0073 & 0.0064 & 1.0000 \end{bmatrix}, t = \begin{bmatrix} 0.0051 \\ 0.0030 \\ 0.0031 \end{bmatrix} \quad (4.3)$$

The rotational and translational RMSE were calculated, equalling 2° and 2.9 mm respectively. The rotational RMSE wasn't as high as the numerical experiments predicted, but the

translational error was within the range of expected, which is too high for the needs of the TUS experiments. Notably, the ICP RMSE output was 1.2 mm and while this seems low in comparison to the limit thresholds, it is not representative of the true error.

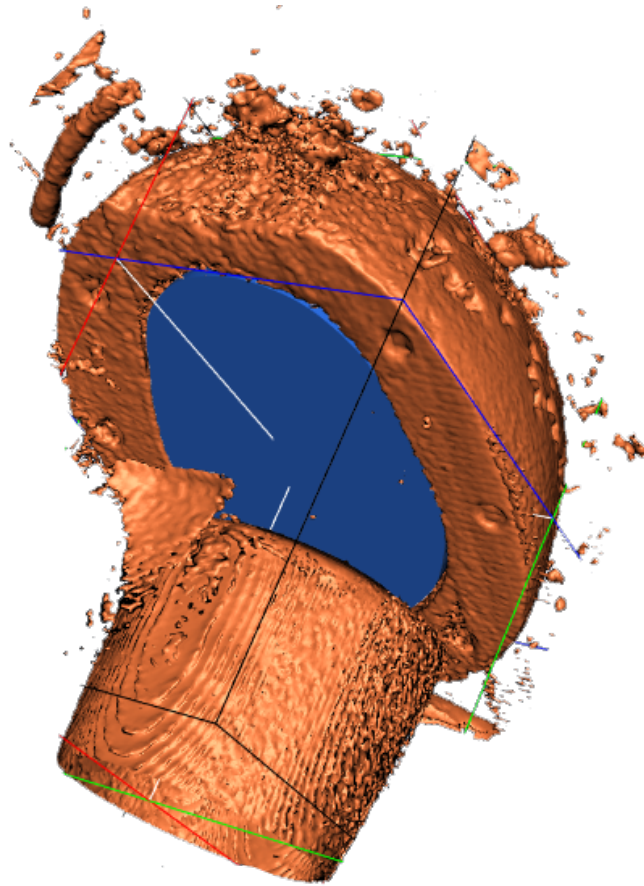


Figure 37: Skull model CAD file overlaid on thresholded MR image to recover \mathbf{T}_2 . The process of overlaying the images was performed by eye.

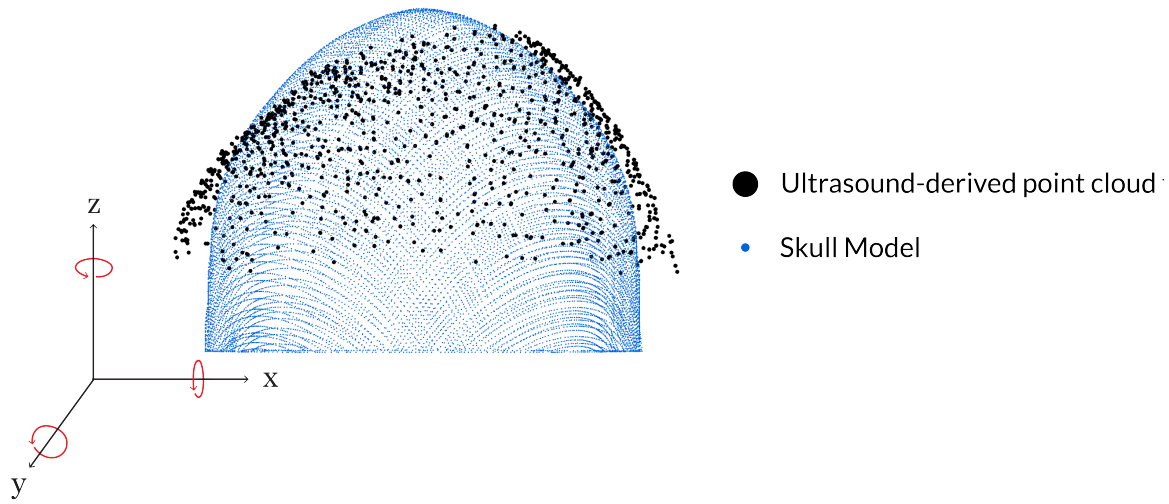


Figure 38: Ultrasound-derived point cloud plotted on top of CAD file-derived point cloud of skull model pre-processing. The ultrasound-derived point cloud looks to be at a 90° angle to the skull model.

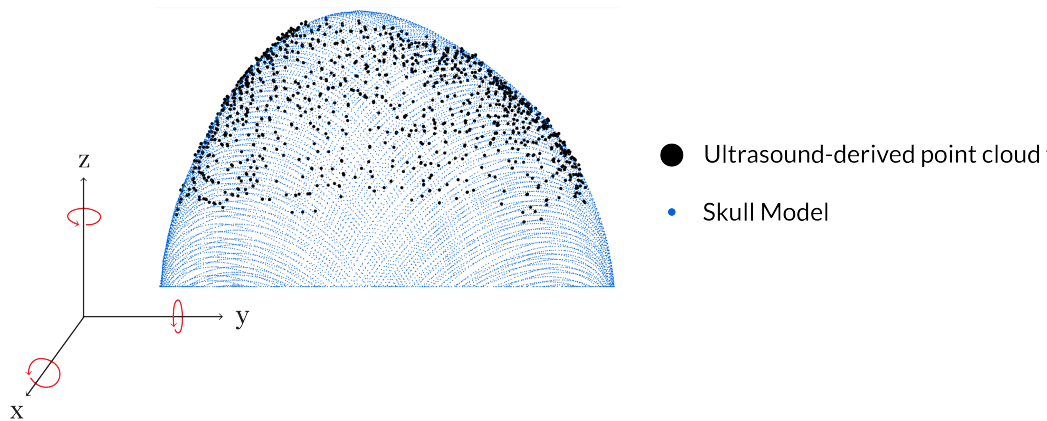


Figure 39: ICP solution plotted on top of CAD file-derived point cloud of skull model. Qualitatively, the solution looks to be converging on the correct solution.

CHAPTER 5

PRELIMINARY EXPERIMENTAL VALIDATION USING HUMAN VOLUNTEER

This chapter describes the preliminary point cloud reconstruction process for a human volunteer and discusses its success. Further tuning of the time of flight picking method is required to choose the reflection from the skull in human subjects.

5.1 Methods

A subject was secured into the TUS helmet using a stereotactic mask to control the location of the subject and to maintain a water tight seal between the skin and the helmet. To maintain a tight seal, the subject should have a shaved head since water can leak out of the holes created by hairs. Unfortunately, in this initial test, the subject's head was not shaved so some water was leaking during the imaging, which resulted in lower or not usable signal from some elements. The subject was taken out of the unit and put back in to try and get a better seal, but the leaking persisted. Regardless, the imaging sequence was performed with varying averages, between 1 average and 32 averages. No pre-processing of the data occurred, so some of the outliers in the point cloud reconstructions were from elements that were not fully submerged.

Additionally, when the unit was taken to Oxford for MR imaging, a subject was placed into the TUS helmet to perform a human validation using an MR image as the ground truth like the skull model. However, the subject also did not have a shaved head and so there was too much leaking to get a proper ultrasound measurement. Further work will be done to create a better water tight seal to transition to human experiments.

5.2 Results

The basic shape of the subject's head can be resolved in the plots in Figure 40. Although, the acquisition and point cloud reconstruction scripts don't pre-process the data before reconstructing a point cloud, so all outliers and unusable signals from elements out of the water can be seen in the images in Figure 40. Qualitatively, the point clouds with more averages seem to have less outliers. Unfortunately, without pre-processing the data it is difficult to tell whether the reduction in outliers is from averaging or better water management. In the future a data classifier function will be implemented prior to point cloud reconstruction.

Another complication noticed in these experiments was that the time of arrival for the human data looked by eye to be inaccurate. The time of flight picking script finds the first time of arrival of the signal. In human experiments, the skin will produce the first reflection, but it is a much smaller reflection than the skull. Figure 41 shows an example set of signals and their calculated time of flights. The skull model has a very sharp clean signal, but the signal from human experiments doesn't have as clear of a start point. This is likely a result of the skin or perhaps air bubbles trapped in the hair. This is an open problem of this method to be worked on in the future, but some ideas include tuning of the time of flight picker, which may result in picking the start of the larger waveform. Also, once the data classifier is implemented, the script should be able to reject signals with little or weak signals like in the fourth line of the human head experiments in Figure 41.

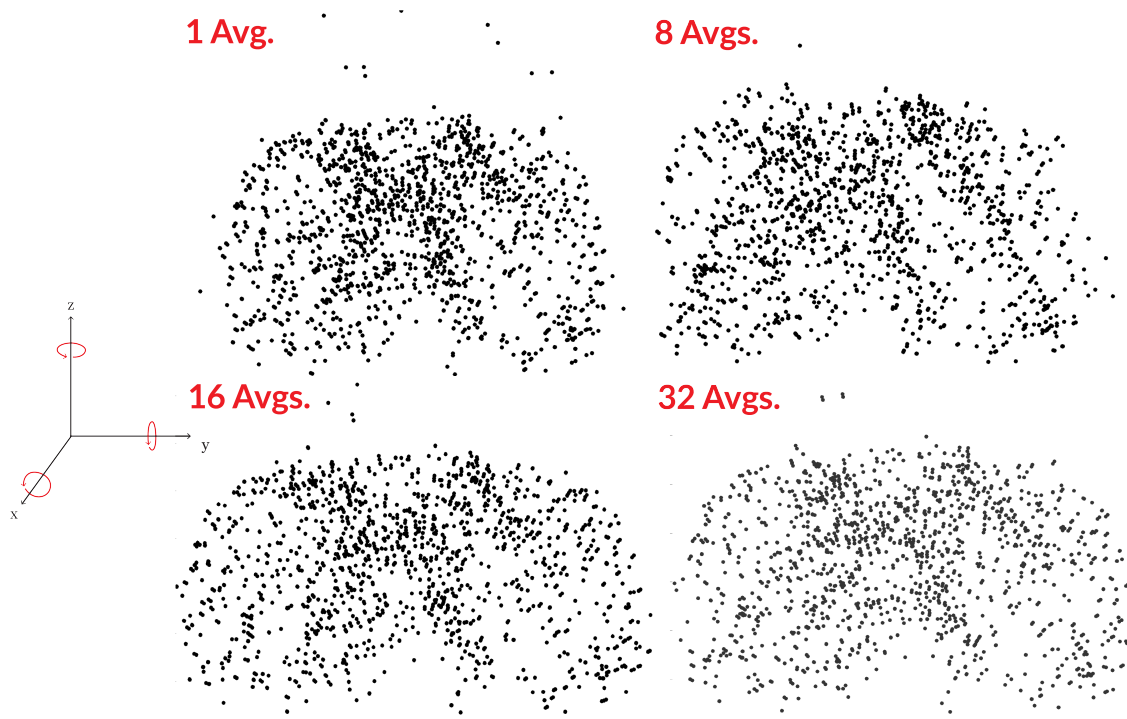


Figure 40: First set of experiments from a subject showing different averaging: 1, 8, 16, and 32 averages. Qualitatively, the image with more averages seem to have less outliers.

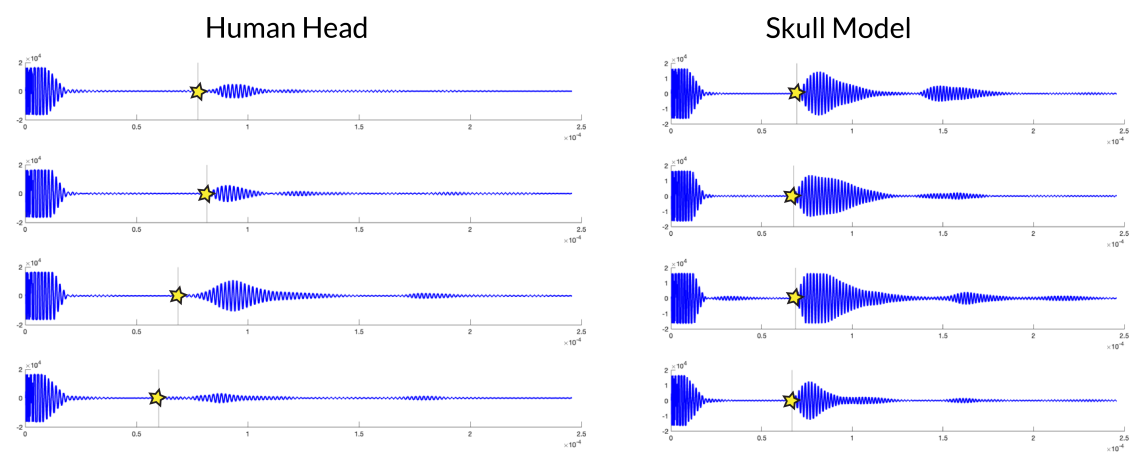


Figure 41: A set of examples of the signals from human and skull model experiments with their associated calculated time of flights marked by the yellow star. This shows that with a human head the time of flights are picked earlier than we would expect to see the skull, likely due to the skin.

CHAPTER 6

CONCLUSION

6.1 Discussion

A method for locating the skull inside the TUS helmet within $\pm 2^\circ$ and ± 2.9 mm was developed and tested numerically and experimentally. The method works by acquiring data from the TUS element array, using reasonable mathematical assumptions to reconstruct a point cloud of the outer surface of the object, and finally registering this point cloud to a pre-therapeutic image-derived point cloud using a multi-start ICP registration technique. It was found numerically that the ICP registration technique is robust to random down-sampling, Gaussian noise, and translational error in the initial starting point. Of note was that with the addition of Gaussian noise, the magnitude of translational and rotational error decreased slightly, from an approximate mean of 1.2° and 4 mm to 0.5° and 2 mm with noise. Adding noise smooths out the error function landscape, making the ICP algorithm less likely to converge on a local minima.

In addition, it was found that the ICP technique was particularly sensitive to any scaling difference between the point clouds, rotations in the initial starting point, and rotational symmetry of the point clouds. As the moving point cloud was scaled, the ICP algorithm solution was increasingly worse, specifically for translation. 20° rotational error and 75 mm translational error were seen. After plotting the solution, it was obvious that the ICP algorithm was shifting the scaled moving point cloud such that there was partial overlap of the point clouds. This error likely comes from the closest point query step of the ICP algorithm and can be corrected for using 1-1 point correspondence or removing the scaling from electrical delays. The algorithm also performed worse when the initial starting point was rotated further from the correct solution. The rotation is by definition a more difficult problem to solve than translation since rotations do not have a large effect on the distance

between the elements and the skull; This distance is the error that the ICP algorithm tries to minimize.

There were additional limitations in the point cloud reconstruction. The TUS elements have a center frequency of 0.5 MHz and are considerably band-limited, so the imaging resolution is limited. However, the time of flight data method finds the first time of arrival of the signal which is unaffected by the pulse length. But, in human experiments the first arrival signal will likely be the skin. The signal from the skin is small, so there is some uncertainty about if the method will pick the skin consistently. The time of flight script will be tuned to pick the start of the larger waveform, which is likely from the skull, in the future.

Using a similar approach to the multi-start ICP method described here, O'Reilly et al. and Crake et al. found registration accuracy of less than 2 mm and 2° which is more favorable than the results from this project, suggesting the registration method can be further improved. With the stereotactic mask, the location of the head will be known within a very small translation and rotation. With such a close initial starting position, the O'Reilly et al. method could be implemented.

6.2 Future Work

Further work is required to upgrade the current version of the ICP algorithm so that it performs as well as the O'Reilly et al. method. Once updated versions are created, validation of the algorithm in experiments that replicate human subject experiments and in human subject experiments are required. Also, further work is needed to tune the reconstruction method to build a point cloud of the skull reflections in human subject experiments.

Preliminary work was completed to validate the algorithm with a better model. Another skull model was designed and printed using a high precision printer and VeroWhite material, which is a material that mimics the acoustic properties of the skull [22]. The skull model can be seen in Figure 42. The skull model was designed such that it has M6 clearance holes

and nut recesses with raised bits to control its location repeatably and accurately. A cross-section of the repeatable locators can be seen in Figure 43. This design was created so that it can be positioned in a testing tank with a custom-designed holder and shifted in precise movements using an electronic 3-axis translation stage. The custom-designed holder is still in the process of being created, but the holder will have contact with the skull with at least three of the repeatable locators, such that all axes of movement will be controlled for. The ICP algorithm can be tested in repeatability experiments. With the electronic translation stage, particular movements can be tested to check the translation numerical experiments in a real world setting. Regrettably, the skull model was not printed in time to perform any experiments, so future work will be done to validate and test the algorithm used for registration using this setup.

For the next version of the ICP algorithm, iterations with added random noise will be implemented to smooth out the error function landscape and the best solution of these iterations will be used, since it was seen numerically that adding noise cuts the error in half. To find the best solution, the ICP algorithm RMSE should not be used since it was shown that the metric was sensitive to added noise. So, the solutions to noise iterations will be used to transform the original, not noisy data and the RMSE will be calculated between the original moving and fixed point clouds. This RMSE will be used to choose the lowest error solution. Also for the next version of the ICP algorithm, an initial starting point from the stereotactic mask will be implemented which will bring the fixed and moving point clouds in close alignment.

Recent work has shown that transcranial ultrasound images can be reconstructed using three-dimensional full-waveform imaging [23]. The method requires a reasonably good initial model of the skull, but with a good initial model, the method can reconstruct images of the brain and skull with sub-millimeter precision. Although computationally expensive, this suggests the possibility of performing an ICP registration using the ultrasound-derived point

cloud and the mesh-derived point cloud to get an initial starting model for the location of the skull. Then, running the full-waveform inversion using the ICP solution as the starting guess. Finally, using the full-waveform inversion solution to calculate the phase shifts induced by the skull. An alternative method could be to perform the ICP registration with the ultrasound-derived point cloud and an average skull and brain model, to resolve the initial model for a full-waveform inversion image. This potential solution would remove the need for a pre-therapeutic image entirely and if successful, also show the possibility of using ultrasound for brain imaging.

There are many other methods for registration, a few with built-in MATLAB functions are phase correlation, normal-distributions transformation, and coherent point drift methods. And, an in-house version of the ICP algorithm can be developed following O'Reilly et al. With the testing function streamlined, these methods can be tested and the data processed to determine which works best for this specific use case.




Figure 42: Skull model 3D-printed out of VeroWhite with visible M6 nut recesses designed and printed to validate the ICP algorithm in experiments that replicate a real world setting.




Figure 43: Cross-section of 3D-printed repeatable locators, designed to hold the skull model in a repeatable and accurate location. The holder holder that these parts fit into will hold the phantom in three spots, so the all possible axes of rotation are controlled for.

REFERENCES

- [1] E. Ottestad and D. S. Orlovich, “History of peripheral nerve stimulation—update for the 21st century,” *Pain Medicine*, vol. 21, no. Supplement_1, S3–S5, 2020.
- [2] K. Hynynen and F. A. Jolesz, “Demonstration of potential noninvasive ultrasound brain therapy through an intact skull,” *Ultrasound in medicine & biology*, vol. 24, no. 2, pp. 275–283, 1998.
- [3] G. T. Clement and K. Hynynen, “A non-invasive method for focusing ultrasound through the human skull,” *Physics in Medicine & Biology*, vol. 47, no. 8, p. 1219, 2002.
- [4] M. A. O’Reilly, R. M. Jones, G. Birman, and K. Hynynen, “Registration of human skull computed tomography data to an ultrasound treatment space using a sparse high frequency ultrasound hemispherical array,” *Medical physics*, vol. 43, no. 9, pp. 5063–5071, 2016.
- [5] W. J. Tyler, S. W. Lani, and G. M. Hwang, “Ultrasonic modulation of neural circuit activity,” *Current opinion in neurobiology*, vol. 50, pp. 222–231, 2018.
- [6] L. Di Biase, E. Falato, and V. Di Lazzaro, “Transcranial focused ultrasound (tfus) and transcranial unfocused ultrasound (tus) neuromodulation: From theoretical principles to stimulation practices,” *Frontiers in neurology*, p. 549, 2019.
- [7] S. Hameroff, M. Trakas, C. Duffield, *et al.*, “Transcranial ultrasound (tus) effects on mental states: A pilot study,” *Brain stimulation*, vol. 6, no. 3, pp. 409–415, 2013.
- [8] W. Legon, T. F. Sato, A. Opitz, *et al.*, “Transcranial focused ultrasound modulates the activity of primary somatosensory cortex in humans,” *Nature neuroscience*, vol. 17, no. 2, pp. 322–329, 2014.
- [9] F. Marquet, Y.-S. Tung, T. Teichert, V. P. Ferrera, and E. E. Konofagou, “Noninvasive, transient and selective blood-brain barrier opening in non-human primates in vivo,” *PloS one*, vol. 6, no. 7, e22598, 2011.
- [10] N. Lipsman, Y. Meng, A. Bethune, *et al.*, *Blood-brain barrier opening in alzheimer’s disease using mr-guided focused ultrasound. nat commun 9 (1): 2336*, 2018.
- [11] C. Jiang, D. Li, F. Xu, Y. Li, C. Liu, and D. Ta, “Numerical evaluation of the influence of skull heterogeneity on transcranial ultrasonic focusing,” *Frontiers in Neuroscience*, vol. 14, p. 317, 2020.
- [12] P. J. Besl and N. D. McKay, “Method for registration of 3-d shapes,” in *Sensor fusion IV: control paradigms and data structures*, International Society for Optics and Photonics, vol. 1611, 1992, pp. 586–606.
- [13] L. Fieten, K. Schmieder, M. Engelhardt, L. Pasalic, K. Radermacher, and S. Heger, “Fast and accurate registration of cranial ct images with a-mode ultrasound,” *International journal of computer assisted radiology and surgery*, vol. 4, no. 3, pp. 225–237, 2009.

- [14] C. Crake, S. T. Brinker, C. M. Coviello, M. S. Livingstone, and N. J. McDannold, “A dual-mode hemispherical sparse array for 3d passive acoustic mapping and skull localization within a clinical mri guided focused ultrasound device,” *Physics in Medicine & Biology*, vol. 63, no. 6, p. 065 008, 2018.
- [15] L. Deng, S. D. Yang, M. A. O’Reilly, R. M. Jones, and K. Hynynen, “An ultrasound-guided hemispherical phased array for microbubble-mediated ultrasound therapy,” *IEEE Transactions on Biomedical Engineering*, 2021.
- [16] N. Lu, T. L. Hall, J. R. Sukovich, *et al.*, “Two-step aberration correction: Application to transcranial histotripsy,” *Physics in Medicine & Biology*, vol. 67, no. 12, p. 125 009, 2022.
- [17] MATLAB, *9.9.0.1524771 (R2020b) Update 2*. Natick, Massachusetts: The MathWorks Inc., 2022.
- [18] U. Castellani and A. Bartoli, “3d shape registration,” in *3D Imaging, Analysis and Applications*, Springer, 2020, pp. 353–411.
- [19] G. G. Slabaugh, “Computing euler angles from a rotation matrix,” *Retrieved on August*, vol. 6, no. 2000, pp. 39–63, 1999.
- [20] C. Rorden, L. Bonilha, J. Fridriksson, B. Bender, and H.-O. Karnath, “Age-specific ct and mri templates for spatial normalization,” *Neuroimage*, vol. 61, no. 4, pp. 957–965, 2012.
- [21] J. Mazziotta, A. Toga, A. Evans, *et al.*, “A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (icbm),” *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, vol. 356, no. 1412, pp. 1293–1322, 2001.
- [22] M. Bakaric, P. Miloro, A. Javaherian, B. T. Cox, B. E. Treeby, and M. D. Brown, “Measurement of the ultrasound attenuation and dispersion in 3d-printed photopolymer materials from 1 to 3.5 mhz,” *The Journal of the Acoustical Society of America*, vol. 150, no. 4, pp. 2798–2805, 2021.
- [23] L. Guasch, O. Calderón Agudo, M.-X. Tang, P. Nachev, and M. Warner, “Full-waveform inversion imaging of the human brain,” *NPJ digital medicine*, vol. 3, no. 1, pp. 1–12, 2020.
- [24] J. Aarnio, G. T. Clement, and K. Hynynen, “A new ultrasound method for determining the acoustic phase shifts caused by the skull bone,” *Ultrasound in medicine & biology*, vol. 31, no. 6, pp. 771–780, 2005.
- [25] L. Deng, A. Hughes, and K. Hynynen, “A noninvasive ultrasound resonance method for detecting skull induced phase shifts may provide a signal for adaptive focusing,” *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 9, pp. 2628–2637, 2020.
- [26] B. K. Horn, “Closed-form solution of absolute orientation using unit quaternions,” *Josa a*, vol. 4, no. 4, pp. 629–642, 1987.
- [27] J. L. Sanguinetti, E. Smith, J. J. Allen, and S. Hameroff, “32 human brain stimulation with transcranial ultrasound,” 2014.

- [28] A. Cheong, *Euclidean distance between two point clouds*, Oct. 2016. [Online]. Available: <https://www.mathworks.com/matlabcentral/fileexchange/59377-euclidean-distance-between-two-point-clouds>.
- [29] Q. Fang, *Iso2mesh*, Sep. 2021. [Online]. Available: <https://github.com/fangq/iso2mesh>.
- [30] A. Javaherian, *Time-of-flight picking each signal*, Dec. 2020.
- [31] A. Javaherian, F. Lucka, and B. T. Cox, "Refraction-corrected ray-based inversion for three-dimensional ultrasound tomography of the breast," *Inverse Problems*, vol. 36, no. 12, p. 125 010, 2020.