Ultra-Wideband Error Modeling for Improved Localization

by

Elizabeth Pedlow

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

Author .....
Department of Mechanical Engineering
August 14, 2021

Certified by .....
John J. Leonard
Samuel C. Collins Professor of Mechanical and Ocean Engineering
Thesis Supervisor

Accepted by .....
Nicholas Hadjiconstantinou
Chairman, Department Committee on Graduate Theses

Ultra-Wideband Error Modeling for Improved Localization

by

Elizabeth Pedlow

Submitted to the Department of Mechanical Engineering
on August 14, 2021, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

Ultra-wideband (UWB) is a modern range measurement technology which can provide high-speed, low-cost ranging, however UWB measurements can be difficult to model. In an effort to increase accuracy of localization using UWB, this thesis develops models to better understand the complex error patterns of UWB range measurements, specifically how separation distance and relative angle between modules affect error. These models are used to develop three error prediction and correction methods to improve localization: (1) range-based error correction, (2) angle-based error correction, and (3) fused range-angle error correction. While it was found that decreasing mean measurement error does not always decrease localization error, the lowest measurement error and lowest localization error both resulted from the fused error correction method. The fused error model combines the separation distance and relative angle models to predict and correct for range error, decreasing the mean measurement error by over 80%, the mean localization error by approximately 35% when using least squares estimation, and by approximately 56% when smoothing the trajectory with a Kalman Filter.

Thesis Supervisor: John J. Leonard
Title: Samuel C. Collins Professor of Mechanical and Ocean Engineering

Acknowledgments

I give my sincerest appreciation to my advisor, Professor John Leonard, whose care about his students surpasses expectation. He took me in as a student even when his lab was full, was infinitely flexible during the unique circumstances of Covid-19, and always provided great insight about whatever interesting data I sent his way.

Thank you to my fellow lab mates in the Marine Robotics Group. Thank you to Nicole Thumma for the time she spent helping me with data collection and for her constant willingness to help. Thank you to Sophia Franklin for her work in building the robots used for the experiments in this thesis and for all of the times she helped me to debug them. Thank you to Antonio Terán for helping to start my research off on the right foot and to Alan Papalia and Kevin Doherty for their technical advice.

I am forever indebted to David and Jeanne-Marie Brookfield. Their generous support during the first semester of my graduate studies allowed me the opportunity of a lifetime. I am truly honored to be the first of a long line of Brookfield Fellows.

Thank you to Professor Sean Humbert of the University of Colorado at Boulder for allowing me the use of his equipment for data collection. His generosity and excitement for everything engineering is truly inspirational.

Thank you to my friends and family who have encouraged me along the way, and especially to Amy Gouws, whose friendship and support know no bounds.

I also owe a tremendous thank you to my fiancé, Leopold Beuken. From discussing technical topics to editing my writing to his daily encouragement, his constant support spanned all aspects of my degree. This thesis would not be complete without him.

I do not have the words to thank my parents, Kevin and Laurie Pedlow, for the time they spent teaching me to enjoy the challenging journey of learning new things. I remember the "backwalkover helper" my dad built me out of two-by-fours to practice gymnastics at home, my first mountain bike ride, racing through homemade "math minutes", and the way my mom could never say no to a new book. All these things encouraged me not only to learn challenging new skills and ideas, but to *love* learning them. The pursuit of my Master's degree started not too long ago, but the journey

began the first day I was born.

I dedicate this thesis to my grandparents, William and Theresa Barna. My grandmother's abilities to design and create are the best of any non-engineer I know, and my grandfather's determination has led him to overcome challenges to which my graduate degree could never compare.

This work was supported in part by the MIT Portugal Program and by ONR MURI grant N00014-19-1-2571 and ONR grant N00014-18-1-2832.

Contents

List of Figures

List of Tables

List of Acronyms

Abbreviation	Meaning
DP	Direct Path
EKF	Extended Kalman Filter
GPS	Global Positioning System
LOS	Line-Of-Sight
NLOS	Non-Line-Of-Sight
RTT	Round Trip Time
TDoA	Time Difference of Arrival
ToA	Time of Arrival
TWR	Two-Way Ranging
UKF	Unscented Kalman Filter
UWB	Ultra-Wideband

Chapter 1

Introduction

Using range measurements to localize an agent's position is referred to as trilateration. Trilateration is a well-researched area and is used by familiar technologies such as the Global Positioning System (GPS). There has been a substantial amount of work on trilateration for robot localization in various domains, such as for underwater vehicles using acoustic beacons, however the problem of localization with Ultra-Wideband (UWB) range measurements remains difficult, due in part to the challenge of modeling UWB range measurement errors. This thesis investigates the problem of localization for small mobile robots using UWB ranging in indoor environments.

UWB has been investigated for several decades, but has only recently begun to see widespread integration into commercial products, such as smartphones and smart home gadgets. It is seen by companies as a strong option for short-distance (typically $<$100 meters), high-speed data transfer (such as Apple's AirDrop feature) and for range-based localization [22]. UWB's wide bandwidth gives its signal the ability to travel through walls and other solid objects, making it a good alternative to traditional ranging technologies (lidar, ultrasound, etc.) for localization in complex environments [21]. In addition to its ability to measure ranges without a direct line-of-sight (LOS) between the modules, UWB is low-power, making it safe around humans and medical devices and a good candidate for use in hospitals [14] and in assisted living facilities [1] for tasks such as patient tracking.

UWB's low-power ranging and ability to function under non-line-of-sight (NLOS)

conditions is advantageous in many scenarios, however UWB still has its challenges. In particular, UWB measurement error can be impacted by many environmental factors and even the distances and angles between the modules. The varying error of UWB range measurements points to the necessity of an error model which can capture the effect of the most influential factors to predict accurately the error of each measurement. The goal of this research is (1) to gain a better understanding of UWB range error and how it changes with separation distance and antenna angle, (2) to develop simple models which capture the essence of these error patterns, and (3) to apply these models to dynamic situations to improve localization accuracy.

1.1 Ultra-Wideband Ranging

Due to its low cost and widespread availability, this work uses the Decawave DWM1001C module which is equipped with Decawave's DW1000 UWB transceiver, running a version of Two-Way-Ranging (TWR) [6, 8]. TWR is based on the principle of multiplying the Time-of-Flight (ToF) of the UWB signal by the speed of light to calculate the distance between the modules.

To measure distances between UWB modules, it is required to have at least two devices: a tag and an anchor. There is no hardware difference between a tag and anchor; the only difference is in the software, which controls what role they each play in communication. The tag is the module attached to the object to be localized which initiates the communication to the surrounding anchor(s). The anchor is typically fixed in place (as will be assumed in this thesis); its job is to listen for a signal from the tag and respond.

There are many different patterns in which communication can happen between an anchor and tag. In particular, Decawave uses a ranging scheme called asymmetric double-sided TWR. The major difference between TWR and asymmetric double-sided TWR is that the former only sends two signals: one from the tag to the anchor and a second from the anchor to the tag. The latter sends three: (1) from the tag to the anchor, (2) from the anchor to the tag, and (3) once more from tag to the anchor.

The advantage of asymmetric double-sided TWR is that it is more robust against clock and frequency drift[1] [6]. A diagram of asymmetric double-sided TWR ranging is shown in Figure 1-1.

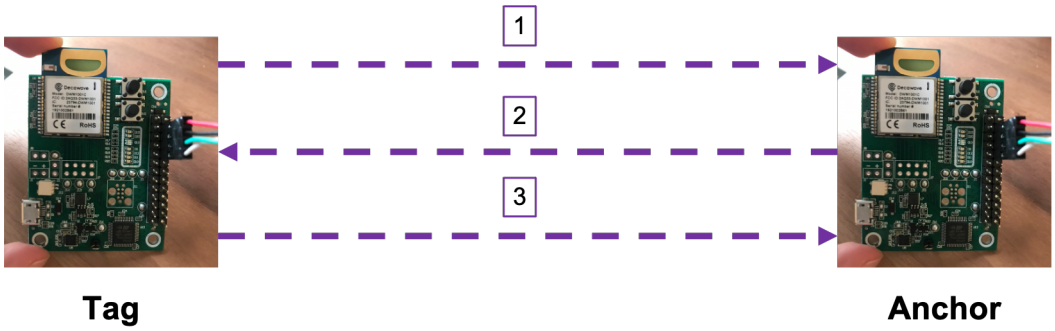

Figure 1-1: TWR uses a three-part signal transmission pattern to determine the distance between the tag and anchor.

From this point forward, asymmetric double-sided TWR will be called TWR for brevity. The TWR equation used to determine the distance between the UWBs is given as

$$T_{prop} = \frac{(T_{1,Tag}T_{2,Anchor}) - (T_{1,Anchor}T_{2,Tag})}{T_{1,Tag} + T_{2,Anchor} + T_{1,Anchor} + T_{2,Tag}} \tag{1.1}$$

where T_{prop} is the estimated time of flight, $T_{1,Tag}$ and $T_{2,Anchor}$ are the respective round trip time times it takes receive the signal back to the module after it was sent. $T_{2,Tag}$ and $T_{1,Anchor}$ are the respective times for the tag and anchor to receive a signal and send one back [6]. A simplified visual representation of these variables is shown in Figure 1-2.

[1]Clock drift is one of the leading causes of UWB range measurement error [24].

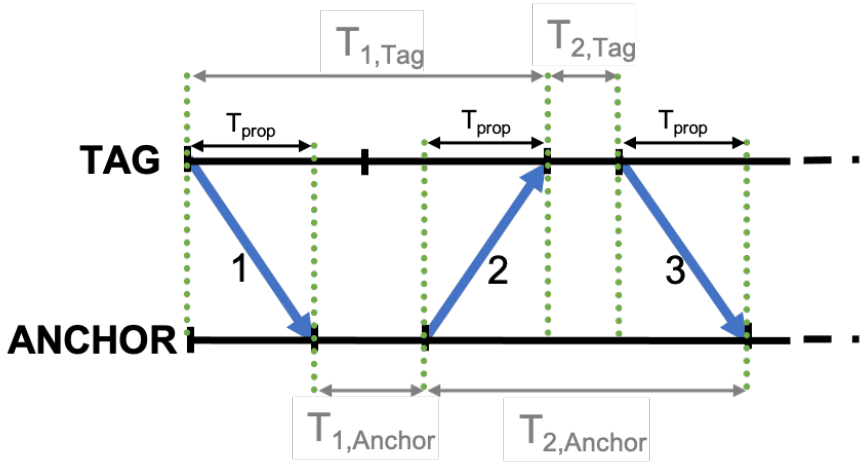

Figure 1-2: A simplified diagram of asymmetric double-sided Two-Way Ranging.

Two of the most researched challenges of UWB are multipath effects and NLOS errors. Multipath effects occur when the signal bounces off objects or walls, often in addition to traveling the direct path (DP) between modules, and therefore can result in *multiple* different distance measurements. One way to mitigate the error from multipath effects is to use the shortest distance returned, since this is usually (but not always) the DP measurement [16]. NLOS errors occur when there is no physical line-of-sight path between the anchor and tag, causing the UWB modules to overestimate the distance between one another. There are two ways this overestimation can happen: the signal can get delayed when it travels through a solid object or, if the object is impenetrable to the signal, the signal may take take a longer path. Both scenarios are shown in Figure 1-3. As long as the geometry of the modules and the object stays the same, the additional length that appears in the range measurement is constant [13], a contributing factor to why NLOS errors are more challenging to detect than multipath errors.

Prorok and Martinoli [21] characterized the error distribution for UWB ranging using Time Difference of Arrival (TDoA) methods. TDoA is similar to TWR, but it requires clock synchronization between modules to measure the difference in arrival times from all anchors to the tag instead of directly measuring the time of flight for each signal separately. They concluded that LOS ranging error can be described by a normal distribution, whereas NLOS ranging error can be described by a log-normal

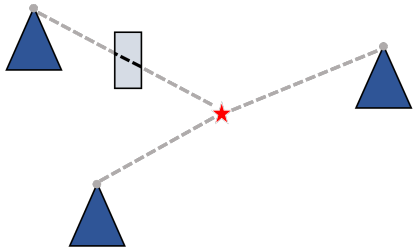

(a) The signal is delayed during its travel through an object.

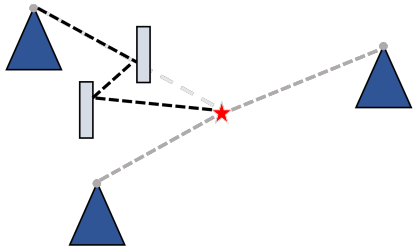

(b) The signal travels a farther distance than the DP.

Figure 1-3: The NLOS path from the beacon to the object is always longer than the LOS path.

distribution. Additionally, their work showed that multiple factors impact the shape of the distribution, including characteristics of the surrounding environment. This means that the parameters describing the measurement error distribution for one environment will not necessarily be valid in another environment.

A growing population of researchers have experimented using machine learning techniques to decrease UWB range and localization error, particularly by mitigating the effects of NLOS ranging [2, 19, 30]. But machine learning has been used for other avenues for error mitigation as well. Tiemann et al. [27] reported significantly decreased range errors when using a neural net to correct for the orientation-dependent portion of error.

In addition to the presence or absence of a LOS path between modules and separation distance, error is also impacted by the angle of the UWB antennas. For example, Sharma et al. [23] found a slight increase in error when antennas were beside each other as opposed to facing each other. UWB antennas are anisotropic in nature, meaning that the module's transmission and receiving of a signal is not uniform in all directions [3, 7, 27]. This is the reason why the error varies with the angle of the

antenna. According to data taken in an anechoic chamber, the orientation most robust against signal strength variation is when the antenna is upright [7], as in Figure 1-4. For that reason, this orientation was used for all the experiments in this thesis.

Figure 1-4: The UWB's upright orientation, used for all experiments.

Ye et al. [31] studied the impact of antenna angle in indoor environments with high levels of multipath effects while using TDoA methods. They found that antenna angle could account for up to a couple centimeters of error regardless of the multipath mitigation algorithm used. Ledergerber and D'Andrea tested incorporating the angle of the UWB antennas into the error model in their localization algorithm to account for the antenna's anisotropic nature. They later abandoned the idea when they saw no significant improvements, however, they did note potential failure points in their work and therefore exploring this should not be written off too quickly [18]. We will consider the effect of antenna angle on measurement error in Chapters 3 and 4.

Despite knowledge of anisotropic antenna signal and error patterns, there is limited information provided by antenna manufacturers that would allow the relationship between antenna angle and error to be inferred [29]. Signal fidelity is one measurement sometimes provided by antenna manufacturers which describes the change in a signal calculated by the correlation between the input and output signals of the antenna [17]. It has been used to classify the three-dimensional performance of antennas [20] and has shown some success when used to predict error due to antenna angle [29].

Cazzorla et al. [3] used Round-Trip-Time (RTT), another UWB distance measure-

ment scheme similar to TWR, to study UWB range error. Under LOS conditions, it was seen that the standard deviation of the range measurements increases as the distance between the tag and anchor increases, in a close to linear fashion. This increase in variance was solely determined from stationary module experiments. A moving tag will further complicate the statistics of the received signal.[2] The effect of distance on error distributions under LOS conditions has also been investigated by De Angelis et al. [5], where the Pearson Type IV distribution was used to model RTT of a UWB signal at various separation distances between the tag and anchor. They found that the distributions became more asymmetric as the tag and anchor became farther apart.

In addition to the geometry between the tag and anchors, electromagnetic interference, humidity, and the operating temperature of the modules are a sample of many environmental factors which can affect range error [5]. At the very least, these factors will vary from one experiment to another, and they may even change during a single experiment or deployment. Additionally, for Decawave's DWM1000 modules, there has been shown to be variation between the modules themselves, causing the measurement error of different anchor-tag pairs to differ [18]. These findings highlight the need to be careful when generalizing any one error model to multiple module pairs.

1.2 Localization

Trilateration is the concept of using distance measurements from multiple known locations to an agent in order to determine its position. Assuming no error and no additional information, a minimum of three range measurements from three stationary beacons is required to determine a point's location in two-dimensional space. With only two range measurements, a pair of possible locations of the point may be determined. With a single range measurement, only a circle on which the point is

[2]The increase in variance of measurements from a moving tag are handled in some NLOS detection algorithms, as noted by Khodjaev et al. [16].

located is known, as shown in Figure 1-5.

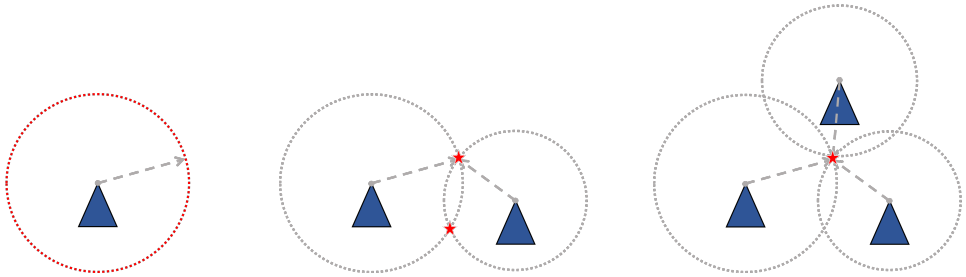

(a) One beacon (b) Two beacons (c) Three beacons

Figure 1-5: Trilateration using one, two, and three beacons.

Whereas Figure 1-5 depicts an idealized scenario, in reality the measured distances will rarely overlap at a single point due to measurement error. Modeling and handling the error of each individual measurement properly is vital to improving the accuracy of the localization estimate.

In general, for range-based localization, least squares estimation is one of the most common algorithms to use. Other solutions include variations of the Kalman Filter, most typically the Extended or Unscented Kalman Filters (EKF and UKF). Through simulations which expressed the simulated range measurements as non-Gaussian distributions (as previously discussed), it was seen that the Particle Filter's state estimation error was lower than both the EKF and least squares estimation [5]. However, Particle Filters can encounter difficulties when the size of the state space grows too large — for example, when performing cooperative localization with many vehicles. The error correction methods that we present later in the thesis can be beneficial regardless of the type of state estimator that is employed.

Sometimes error in the localization estimate from these algorithms can be induced by an assumption that ties together range measurements taken at slightly different times. This becomes a problem when the measurement rate and robot's speed are such that the robot's position changes significantly between measurements from different anchors. One approach to account for this issue is to consider both the time and distance of each range measurement to fit a trajectory to the data [10, 11, 25, 26]. While slower robots that stay close to the anchors may be able to sample at high enough

rates that this issue becomes negligible, Clark et al. [4] suggest that localization algorithms applied to faster robots or with limited sampling rates may be hindered by time delays. On the other hand, it is also important to note that position estimation can also fail when the sampling rate is too high, since this can augment the size of the estimation problem such that it is too large to feasibly solve [11]. From this point forward, it will be assumed that the sampling rate of the UWB modules and the speed of the robots are such that the distance the robot travels between samples is negligible and that the frequency of measurements results in a problem which is not too large to solve.

Using a UKF, Fu et al. [9] concluded a notable pattern to localization error using Time of Arrival (ToA) range calculations with UWB modules. They found that that the error of the two-dimensional position estimate follows a radial pattern. More specifically, the direction of the error points away from the geometric center of the square pattern of anchors, and the error magnitude is correlated with the distance from that point. To compensate for this error, changed the UKF to adjust for the expected error as a function of the distance and direction of the estimated location from the geometric center.

While the work by Fu et al. [9] shows that it is possible to compensate for the radially-patterned error post-localization calculation, this solution does not seem to target the origin of the error. The radial error pattern suggests that measurement error may be distance-dependent. The following work in this thesis is focused on the distance- and angle-dependent error of UWB range measurements and develops models to predict that error with the ultimate goal of improving localization accuracy.

Chapter 2 describes the robot used for testing and presents a model for the robot's kinematics. It also describes the mathematics of range-based localization using least squares estimation. Chapter 3 seeks to understand patterns of UWB range error for various separation distances and antenna angles using static data. Subsequently, Chapter 4 applies the conclusions from the static data to develop and compare three different models for predicting and correcting range measurement error. These models are then evaluated for the potential to improve localization performance.

Chapter 2

Robot Localization

The robot used for all experiments was designed and built by the Marine Robotics Group, the research group of which the author is involved. The robot is a differential drive robot, meaning that the robot has one bi-directional wheel on either side and one low-friction "peg" which slides passively on the ground. The robot was designed to be low-cost, yet reliable. An image of the robot is shown in Figure 2-1.

Figure 2-1: The robot used for all experiments.

2.1 Development of the Robot's Equations of Motion

The state vector is defined as the robot's two-dimensional pose (x, y, and θ) with respect to a global frame, as shown in Figure 2-2.

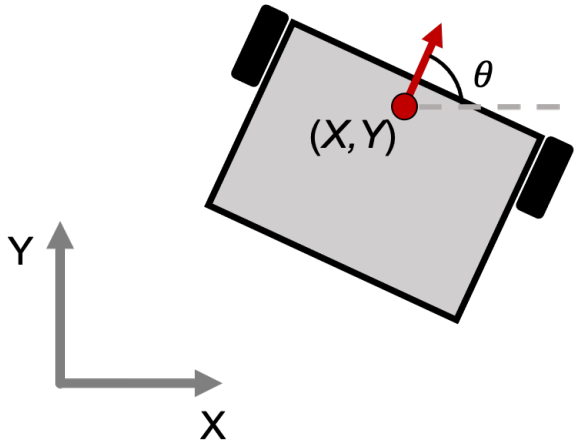

Figure 2-2: The top-view of the robot, showing its pose with respect to the global frame. The red circle indicates the point on the robot which is used to define its location.

To model the differential drive robot, the unicycle model was used. The variables X and Y are used to indicate the global frame, and the variables x and y are used to indicate the body frame. The two inputs to the system are the desired left and right wheel velocities, however the discrete state space model is set up so that the input vector is the change in X, Y, and θ for the current time step. The final propagation equation is written as

$$\begin{bmatrix} X_{t+1} \\ Y_{t+1} \\ \theta_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X_t \\ Y_t \\ \theta_t \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} dX_t \\ dY_t \\ d\theta_t \end{bmatrix} \tag{2.1}$$

where the subscript t denotes the previous time step and $t+1$ denotes the current time step. The inputs to the matrix dX_t, dY_t, and $d\theta_t$ are the changes in X, Y, and θ from the previous time step. They are calculated by

$$\begin{aligned} dX_t &= dx_t \cos\theta - dy_t \sin\theta \\ dY_t &= dx_t \sin\theta + dy_t \cos\theta \\ d\theta_t &= 2 \arcsin\left(\frac{dr_{r,t} - dr_{l,t}}{2b}\right) \end{aligned} \tag{2.2}$$

where b is the wheelbase, $dr_{r,t}$ and $dr_{l,t}$ are the linear distance traveled by the right and left wheels, and dx_t and dy_t are the changes in the x and y directions of the body coordinate frame. The inputs to Equations 2.2 are calculated by

$$\begin{aligned} dr_{r,t} &= R_w u_{r,t} dt \\ dr_{l,t} &= R_w u_{l,t} dt \\ dr_t &= \frac{1}{2}(dr_{r,t} + dr_{l,t}) \\ dx_t &= dr_t \cos\left(\frac{d\theta_t}{2}\right) \\ dy_t &= dr_t \sin\left(\frac{d\theta_t}{2}\right) \end{aligned} \tag{2.3}$$

where R_w is the wheel radius, $d_{r,t}$ is the distance traveled by the robot, $u_{r,t}$ is the velocity input to the right wheel, and $u_{l,t}$ is the velocity input to the left wheel. The two constants in this model, b and R_w, are visually represented in Figure 2-3.

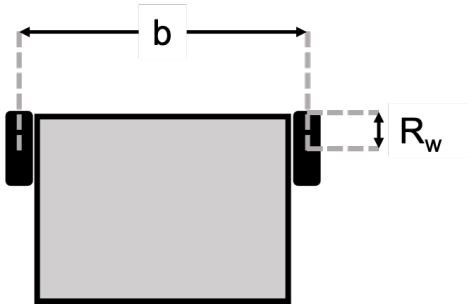

Figure 2-3: The measurements of the wheelbase, b, and the wheel radius, R_w, on a top-view diagram of the robot.

To test the validity of this model, consider a robot traveling a lawnmower pattern.[1] Eight trials were analyzed for model validation. A single trial was used to generate all of following plots since its mean position and heading errors were most similar to the mean errors averaged over all trials.

The starting position and heading for state propagation are equivalent to the starting point of the ground truth measurements. A plot of the ground truth data taken by a Vicon motion capture system with the propagated state is shown for the translational movement in Figure 2-4.

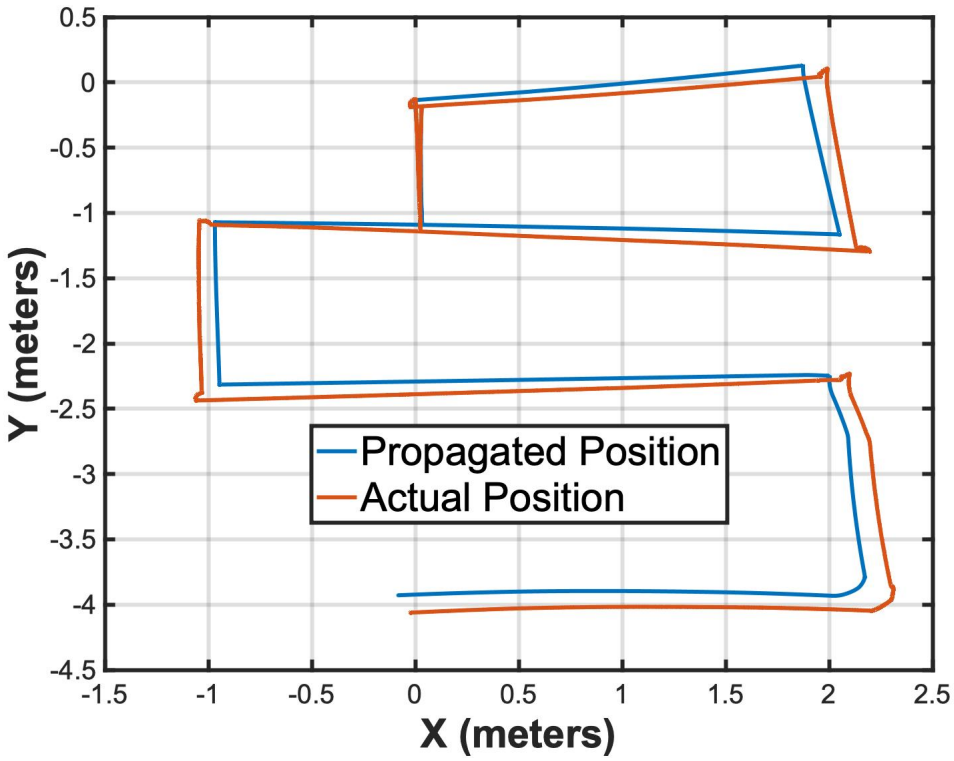

Figure 2-4: Propagated and ground truth positions (from a motion capture system) of the robot's lawnmower trajectory.

[1]The same experiments are used for model validation as were used to test the three error correction methods developed and analyzed in Chapter 4.

The individual errors in the X and Y directions are shown in Figures 2-5a and 2-5b. Averaged for all trials, the mean absolute value error in the X and Y directions are 7.1 and 5.4 cm, respectively. This translates to a mean position error of 10.0 cm throughout the two-and-a-half minute experiments and a final position which has, on average, an error equivalent to 0.9% of the total distance traveled. Figure 2-5c shows the heading error. The mean absolute value of heading error for all trials was 1.6 degrees.

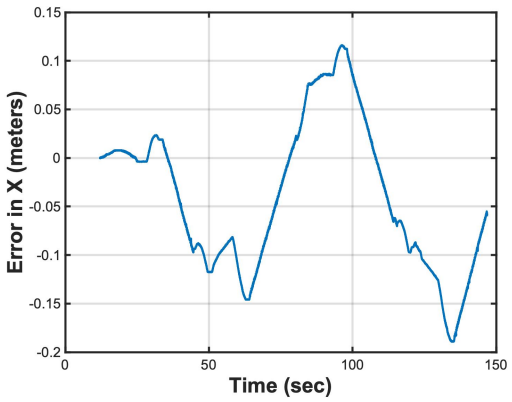

(a) Error of X position

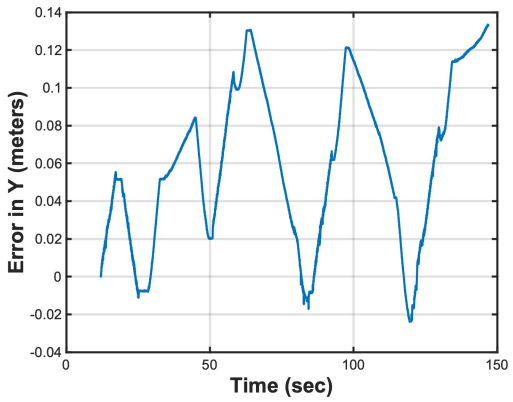

(b) Error of Y position

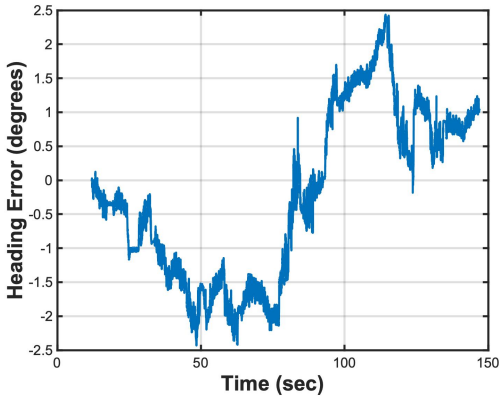

(c) Heading Error

Figure 2-5: The robot's position and heading were estimated by propagating the equations of motion. These figures show the error in the X, Y, and θ directions of the estimate as compared with ground truth measured by a motion capture system.

Overall, the translational motion using the propagated equations slightly under-estimated the actual distance traveled. The estimated heading of pure rotation using

the propagated equations had high accuracy, with the maximum mean heading error of any single trial only 2.6 degrees. However, it must be noted that due to wheel slip, the accuracy of this model can vary considerably depending on the ground surface and the length of the trajectory.

2.2 Localization using Least Squares Estimation

Since the range measurements are noisy, three range measurements will likely never align perfectly at a single point. To estimate the position of the robot using noisy measurements, least squares estimation was used.

The measurement equation was assumed to be of the form

$$\mathbf{p} = \boldsymbol{\phi}^T \hat{\mathbf{X}} \tag{2.4}$$

where $\hat{\mathbf{X}}$ is a 2x1 vector of the estimated X and Y coordinates of the robot in the global frame, $\boldsymbol{\phi}^T$ is a 3x2 matrix, and $\mathbf{p}$ is a 3x1 matrix which is recalculated at each time step. The matrices $\boldsymbol{\phi}^T$ and $\mathbf{p}$ are defined as

$$\boldsymbol{\phi}^T = \begin{bmatrix} 2(X_2 - X_1) & 2(Y_2 - Y_1) \\ 2(X_3 - X_2) & 2(Y_3 - Y_2) \\ 2(X_1 - X_3) & 2(Y_1 - Y_3) \end{bmatrix} \tag{2.5}$$

and

$$\mathbf{p} = \begin{bmatrix} (r_1^2 - X_1^2 - Y_1^2) - (r_2^2 - X_2^2 - Y_2^2) \\ (r_2^2 - X_2^2 - Y_2^2) - (r_3^2 - X_3^2 - Y_3^2) \\ (r_3^2 - X_3^2 - Y_3^2) - (r_1^2 - X_1^2 - Y_1^2) \end{bmatrix} \tag{2.6}$$

where (X_1, Y_1), (X_2, Y_2), and (X_3, Y_3) are the positions and r_1, r_2, and r_3 are the range measurements of beacons 1, 2, and 3, respectively.

This least squares estimate is used in Section 4.5 to estimate the robot's position

using measurements corrected by the three different error models developed in Sections 4.2, 4.3, and 4.4. Subsequently, the estimated positions from the least squares algorithm are combined with the equations of motion developed in Section 2.1 using a Kalman Filter. The results from the least squares estimate and from the Kalman Filter are used to compare the localization accuracy of the three error correction methods.

Chapter 3

Static Error Modeling

3.1 Separation Distance Data

The most basic data which can give insight into patterns of error distribution within UWB ranging is data taken while the tag and the anchor are both fixed in place. As a starting point for understanding the impact of the separation distance of modules on measurement error and to determine how those trends differed for LOS and NLOS paths, static ranging data were collected in two different indoor spaces. One environment was a large, open motion capture space and the other was a home with a reasonably open floor plan and uncluttered environment.

During data collection in each location, the tag was kept in one location throughout the experiment, while the anchors were moved to varying distances between one and eleven meters from the tag. The names of the anchors are consistent with the id's given to each UWB module by Decawave: *dw9838*, *dw971f*, and *dw2f95*. At each distance, the anchors were fixed in place while roughly 5000 measurements were collected between each anchor-tag pair. At all times during data collection, the tag and anchors were oriented vertically and facing each other, as depicted in Figure 3-1. It is important to note that the exact angles of the UWB modules were not considered since it was assumed that minuscule differences in relative angles between the tag and each anchor were negligible. Environments were kept as constant as possible, with the only deviations being minimal human movement far from both the tag and sensor. It

was assumed that this slight motion would have negligible impact on the UWB range measurements. Ground truth data in the motion capture space was determined by Vicon cameras and was hand-measured in the home environment, resulting in much lower uncertainty in the motion capture space. For this reason, the results from the laboratory data set were trusted more when drawing conclusions about UWB error characteristics.

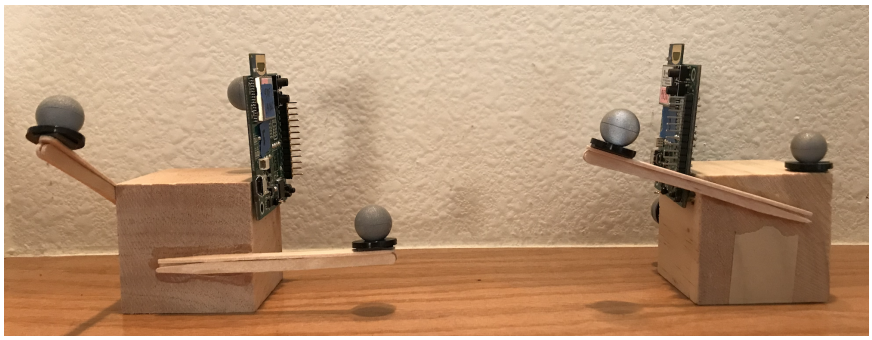

Figure 3-1: UWB modules facing each other. The blocks on which they are mounted hold ultra-reflective spheres used to track the objects in the motion capture space.

Figure 3-2 presents the data from the static ranging experiments taken in the motion capture space. Figures 3-2a and 3-2c show that the average error magnitude increases as separation distance increases in a close to linear trend for both LOS and NLOS ranging. Additionally, it can be seen that the error differs for each anchor-tag pair. This is unlikely to be caused by environmental factors since the anchors were positioned within centimeters of each other and were ranging to the same tag at the same time.

Further analysis of Figure 3-2 reveals a key feature of the data: smaller separation distances tend to vary more in their expected error. More specifically, around 3000 mm separation distance, there is an abrupt change in the variance of the measurements from higher-variance (seemingly more "random") to lower-variance with a more predictable trend. To better understand this trend, a line can be fit to the data above 3000 mm separation distance using least squares estimation. The separation distance,

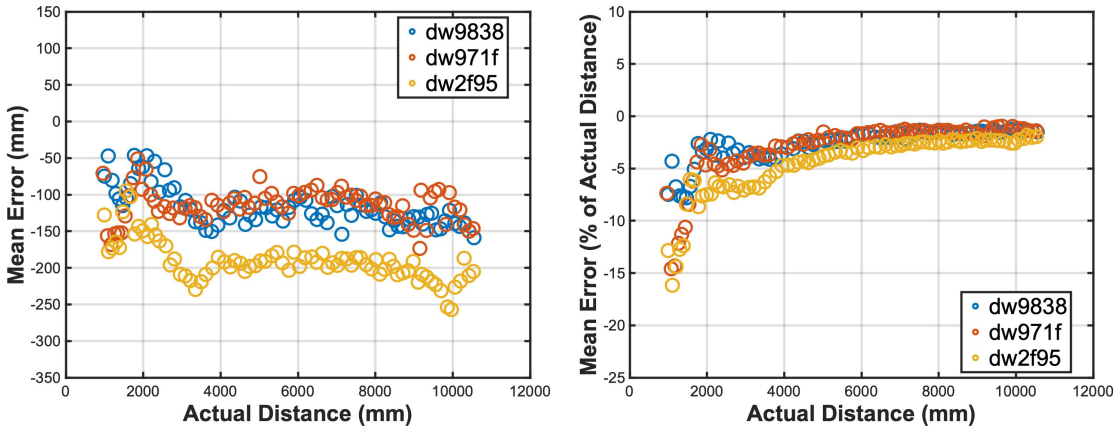

(a) LOS error

(b) LOS percent error

(c) NLOS error

(d) NLOS percent error

Figure 3-2: LOS and NLOS range error in the motion capture space.

R, between the anchor and tag is defined as

$$R = \hat{R} - E \tag{3.1}$$

where $\hat{R}$ is the measured range between the UWBs and E is the error. The linear relationship between separation distance, R, and error, E, is defined as

$$E = mR + b \tag{3.2}$$

where m and b are the constants found for each anchor-tag pair. The values of these constants are shown in Table 3.1. The lines of best fit are shown alongside the raw

data in Figure 3-3. While the parameter b of the best fit lines may differ for various anchor-tag pairs in a given environment, the rate at which the error changes as a function of separation distance, m, is consistent across all anchor-tag pairs. In fact, this trend is so consistent that if a line is fit to all data points above 3000 mm using least squares estimation, the slopes for all three anchor-tag pairs are similar, with a maximum percent difference of 16%.

Table 3.1: Distance model parameters.

Anchor	m [1]	b [mm]
dw9838	-1.93×10^{-3}	-115.4
dw971f	-1.93×10^{-3}	-101.4
dw2f95	-2.24×10^{-3}	-186.6

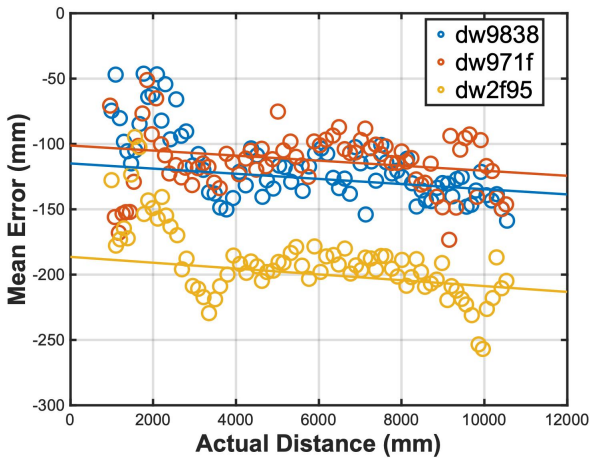

Figure 3-3: LOS error in the motion capture space with lines fitted by least squares estimation. The lines are described by Equation 3.2 and the parameters for each line are included in Table 3.1.

Though analyzing mean error values provides a solid understanding of general error trends, it does not completely describe the error *distributions*. To better understand these data, standard deviation was considered. It was seen that for LOS and NLOS measurements in both environments, the standard deviation increased in a linear fashion with respect to separation distance. This pattern is shown for both

locations and paths in Figure 3-4. It can also be seen that the standard deviation of LOS and NLOS measurements taken in the motion capture space both include a section of data which stands out among the rest, approximately 8000-9500 mm for the LOS data and approximately 5000-9000 for the NLOS data. This outlying trend was consistent for both LOS and NLOS data despite being taken days apart, but did not occur when ranging with the same UWBs in the home environment. Even though the distances at which these outlying data occur are not exactly the same in the LOS and NLOS data, it is hypothesized that this is the result of a physical characteristic of the motion capture space. This unique data pattern is one example of the difficult-to-model nature of UWB ranging which has led researchers to varying conclusions regarding UWB error distributions.

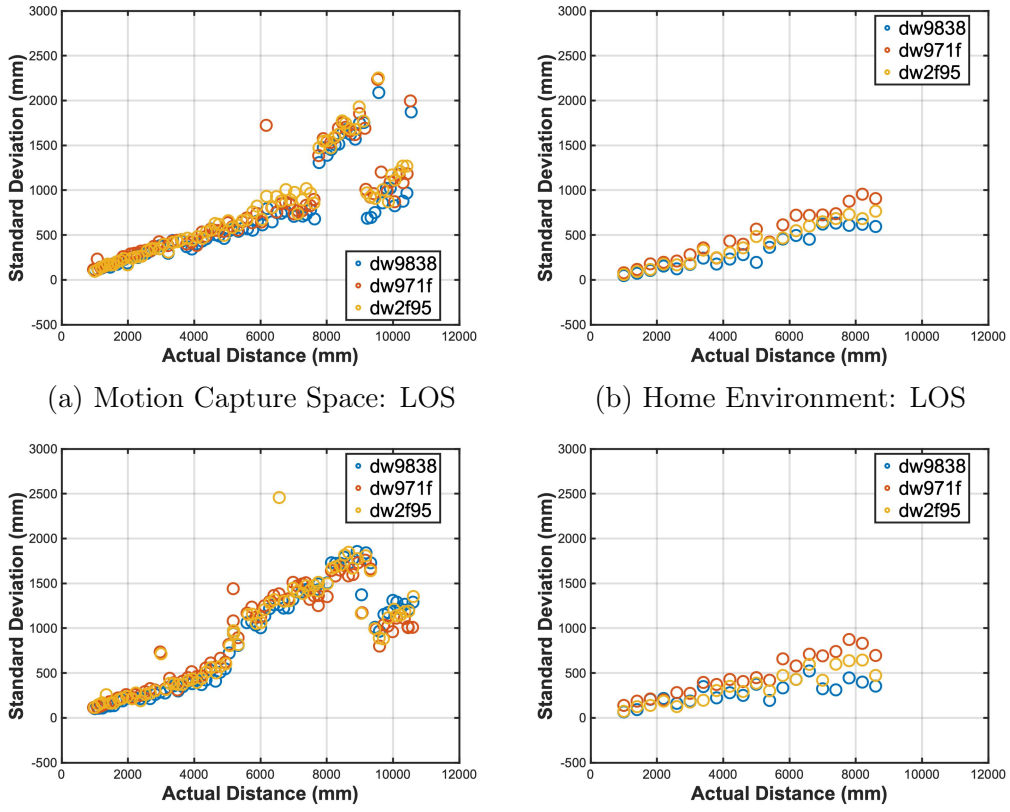

(a) Motion Capture Space: LOS

(b) Home Environment: LOS

(c) Motion Capture Space: NLOS

(d) Home Environment: NLOS

Figure 3-4: The standard deviation of the stationary range measurements at different distances for LOS and NLOS paths in two locations.

3.2 Relative Angle Data

Section 3.1 showed that there is a clear correlation between separation distance and measurement error, although additional data taken with a moving tag and static anchors showed that there is likely more to the story. In data collected where a tag was mounted on a robot which drove a lawnmower pattern in the motion capture space, it was seen that error of the estimated position corresponded to the direction of the robot's travel. More specifically, the magnitude of the position error was higher when the robot was traveling parallel to the y-axis rather than parallel to the x-axis, as seen in Figure 3-5.[1] These results suggest a correlation between sensor angle and range error. To quickly test this hypothesis, the tag mounted on the robot was rotated by 90 degrees clockwise and a similar lawnmower path was driven. This time, the higher magnitude error was seen while the robot was driving parallel to the x-axis, as seen in Figure 3-6. This analysis confirms the directional nature of the UWB antennas. To further investigate this behavior, additional static data were taken where the distance between the anchor and tag remained constant, but the angle of the tag changed.

The experiment was conducted in the motion capture space, with ground truth data for position and orientation collected by the Vicon camera system. A diagram of the experimental setup is shown in Figure 3-7. All three anchors were positioned beside each other, and the distance to the tag was approximately 7400-7500 mm, a distance purposefully chosen to be greater than 3000 mm — the distance around which the error in the separation distance data became significantly more predictable. As with the previous static data set, this static angular data also collected around 5000 range measurements from each of the three anchors while all UWB modules were fixed in place. The tag was rotated so that data were captured roughly every 9 degrees, or at 40 different points within a 360 degree circle.

[1] More details about this experimental setup and an analysis of the results are presented in Chapter 4.

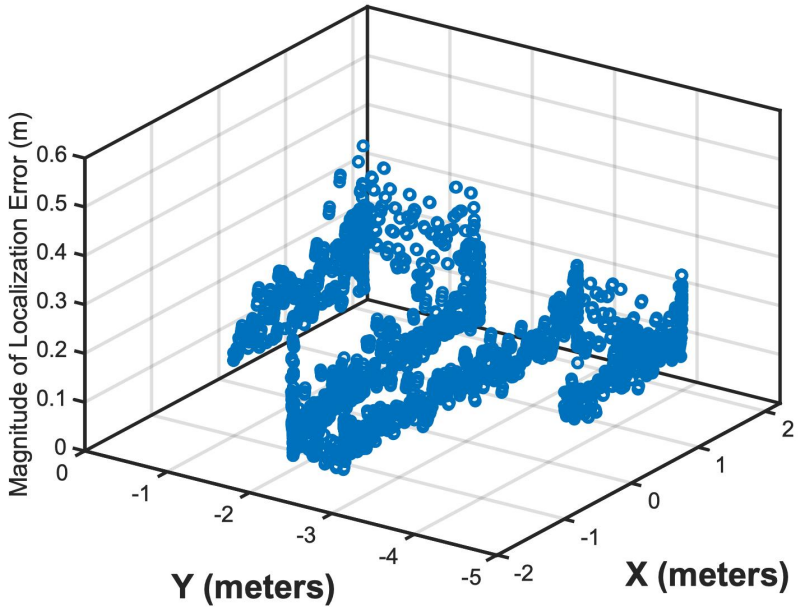

Figure 3-5: A three-dimensional plot of the magnitude of localization error of the robot driving a lawnmower pattern with the antenna pointing straight ahead on the robot.

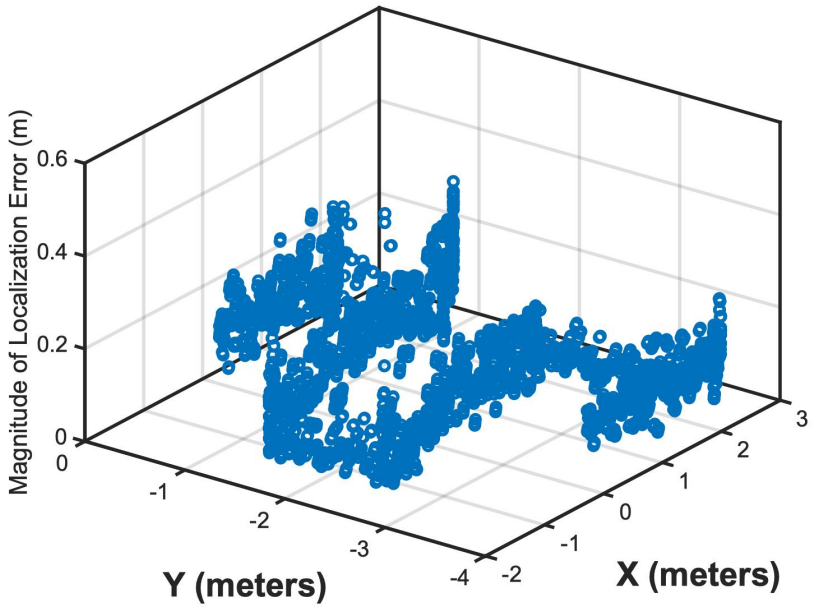

Figure 3-6: A three-dimensional plot of the magnitude of localization error of the robot driving a lawnmower pattern with the antenna rotated 90 degrees clockwise.

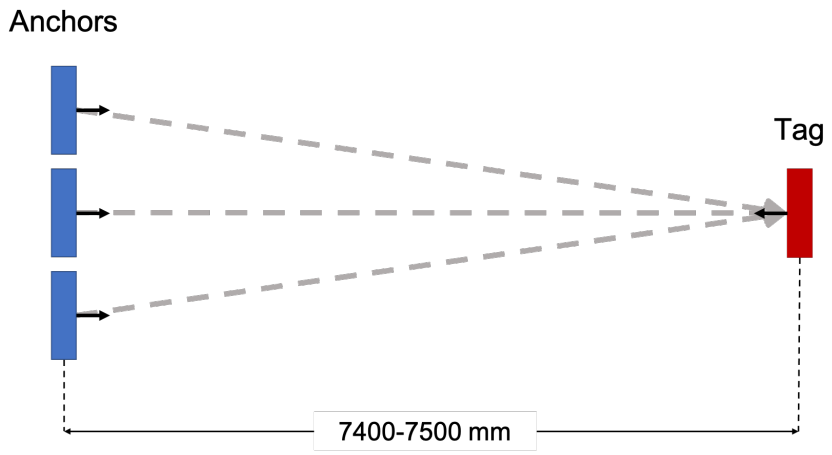

Figure 3-7: Angular data experimental setup.

Since the geometry of the setup causes the angles between each anchor and the tag to be slightly different, a metric for the *relative angle* was designed to normalize for this difference. The relative angle was defined as the difference between the tag's heading in the global coordinate frame and that of the anchor such that the relative angle, θ_{rel}, is calculated by

$$\theta_{rel} = \beta - \alpha \tag{3.3}$$

where β is the heading of the tag and α is the heading of the anchor, as shown in Figure 3-8. For reference, an anchor and tag facing each other would have a relative angle of 180 degrees.

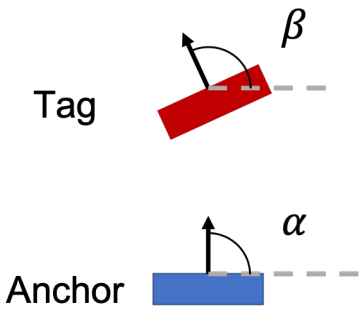

Figure 3-8: A visual representation of heading parameters for Equation 3.3.

The error magnitude versus relative angle is plotted for each anchor-tag pair in

Figure 3-9. It is important to point out that it is the *magnitude* of the mean percent error plotted radially; the actual percent errors were all negative, meaning that range measurements, on average, underestimated the actual distance. The data seem qualitatively elliptical in shape, therefore an open-source ellipse-fitting MATLAB function was used to model the data [12]. The elliptical models are overlaid in red on the raw data for each anchor-tag pair in Figure 3-9. It can be seen that the lowest error occurs with a relative angle of 0 degrees. Although this is consistent across the data for all three anchors, it is not explicitly reflected in the ellipse model for simplicity and not to overfit the data. Future work could create a more detailed angular error model.

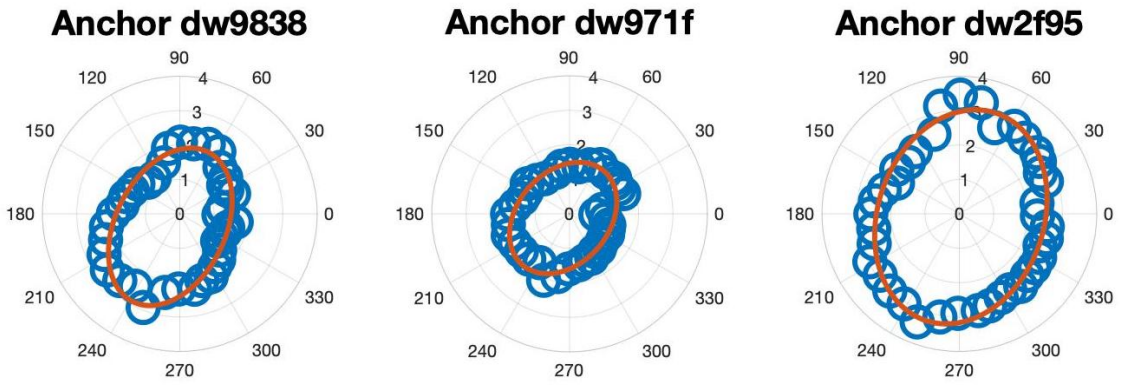

Figure 3-9: The static relative angle data taken in the motion capture space, with the overlaid elliptical models. The radius represents the magnitude of percent error for each relative angle.

To understand the impact of relative angle on range error, consider a scenario where relative angle has no impact on range error. In this case, the static angle data points would form a perfect circle, or a nearly-perfect circle when including unavoidable sensor noise. To numerically confirm the significant difference between a circle and the resulting ellipses, eccentricity can be calculated. The eccentricity e of an ellipse can assume any value $0 \leq e < 1$. A circle is a special case of an ellipse where $e = 0$ and a parabola's eccentricity is $e = 1$. For anchors *dw9838*, *dw971f*, and *dw2f95*, the eccentricities of the elliptical error models are 0.7547, 0.6748, and 0.6510, respectively. These calculated eccentricities of the angular models show that

a perfect circle would not suitably model the data.

Parallels between the static distance data set (Section 3.1) and this static angular data set highlight the consistency of the UWB ranging error in a single environment — even when data is taken several weeks apart. The relative angle data were taken at a separation distance of approximately 7400-7500 mm, and the separation distance data were taken with a relative angle of 180 for all distances. Comparing the angular model's percent error estimate for a relative angle of 180 degrees and the separation distance model's percent error estimate at 7450 mm shows that the two estimates coincide well. Of all anchor-tag pairs, the maximum difference between the estimates is approximately 11%. The values for each anchor-tag pair can be seen in Table 3.2. While this does not prove the models accurate, it is evidence of the models' validity and of the consistency and predictability of the error of an anchor-tag pair — at least in a given environment. This evidence suggests that a single error model for one environment will be sufficient to decrease effectively the range measurement errors for the entire deployment.

Table 3.2: Measured and predicted percent errors.

	Model-Predicted Percent Error (%)		Percent
Anchor	Separation Distance	Relative Angle	Difference
dw9838	-1.86	-1.74	6.9%
dw971f	-1.64	-1.55	5.8%
dw2f95	-2.42	-2.73	11.4%

3.3 Chapter Summary

This chapter has presented static data which explore the impact of separation distance and relative angle on range measurement error. It was found that the magnitude of error increased with separation distance and that the error changed based on the relative angle between the anchor and tag. Two models were developed to described the error patterns. First, a linear model was used to describe the relationship between

separation distance and error. Second, an elliptical model was used to describe the magnitude of percent error as a function of relative angle.

In the next section, these two error models are applied to the problem of robot localization to correct for error in range measurements. Three methods of error prediction are developed and analyzed. The first relies only on the linear separation distance error model, and the second relies only on the elliptical relative angle error model. Finally, the *fused model* combines the two former models for a more thorough representation of the error patterns. Each of the three methods is analyzed for its ability to decrease measurement and localization error.

Chapter 4

Correcting UWB Measurement Error

Chapter 3 modeled UWB measurement error through static ranging experiments. This chapter extrapolates those findings to the localization of a moving robot in order to predict and correct for error of individual range measurements. By decreasing range measurement error, the accuracy of robot localization can be improved.

Section 4.1 describes the experimental setup for the data used to test error correction methods. The subsequent sections develop three methods to predict and correct for range measurement error. Section 4.2 uses a linear model to correct for error due to separation distance. Section 4.3 uses an elliptical model to correct for error due to the relative angle between the antennas. Finally, Section 4.4 fuses the separation distance and relative angle error models to develop a new model which is used to correct for range measurement error. The three models are compared for their ability to decrease range error and to improve localization accuracy in Section 4.5.

4.1 Experimental Setup

There are two sets of data, mentioned briefly in Section 3.2, which were used to validate the effectiveness of the error correction methods. Both data sets were taken in the motion capture space and consisted of a robot carrying a UWB tag while ranging to three anchors fixed in space. The anchors were purposefully placed such that the separation distance was always greater than 3000 mm — the distance above which

ranging error becomes more predictable. The robot's path began with a *calibration period* in which the robot drove roughly one meter straight forward, reversed back to its starting position, and then rotated counterclockwise by 90 degrees to begin a lawnmower pattern. An example ground truth path from the data is shown in Figure 4-1. It is important to note that there was variation in the calibration period and lawnmower pattern of each trial since the robot was controlled by a joystick.

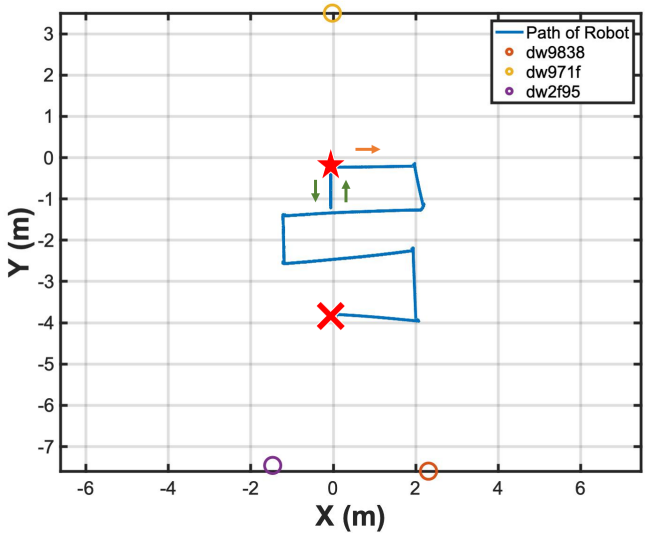

Figure 4-1: An example lawnmower path. The red star and X indicate the robot's start and finish points. The green arrows show the robot's direction of travel during the calibration period, and the orange arrow shows the robot's direction of travel at the start of the lawnmower pattern.

The tag in the front-facing antenna data set was mounted such that the front of the antenna was parallel to the front of the robot. The right-facing antenna data set had the tag rotated 90 degrees clockwise, or facing right on the robot. Photos of the robot and tag setup for each data set are shown in Figure 4-2.

It was a purposeful design choice to build the calibration period into the experiments instead of simply relying on the exact error models developed in Chapter 3. As discussed in Chapter 1, there are a variety of environmental factors, such as temperature, humidity, and electromagnetic interference, of which we had no control. By using a calibration period, the error models could be adjusted to the current environmental conditions, possibly providing higher accuracy error correction.

(a) Tag antenna aligned with the front of the robot to collect the front-facing antenna data set.

(b) Tag antenna 90 degrees offset from the front of the robot to collect the right-facing antenna data set. (Antenna is facing to the left in the photo.)

Figure 4-2: Orientation of the tag's antenna on the robot for two separate data sets.

The parameters of the linear and elliptical models from Chapter 3 were averaged across all three anchors to develop two *generalized* models: one for separation distance and one for relative angle. All three error correction methods began by assigning the generalized linear and/or elliptical model(s) to each of the three anchors. Then, the models were uniquely adjusted to better fit the error patterns of each anchor by using data from the calibration period. This was done by propagating the state equations from Section 2.1 and using them as ground truth to calculate estimates for the measurement error of each anchor. The most important point in all of this is that it means these error correction methods do not require the exact static error models to be determined for each anchor prior to deployment.

Before continuing, consider the validity of using the propagated state equations as ground truth during this calibration period. It was seen that the mean translational error during the calibration period was 26 mm. The calibration period has a minimum distance of 3500 mm to all anchors. Therefore, this mean error translates to a *maximum* mean range error of less than 1%.

4.2 Error Correction using Separation Distance

It was seen in Section 3.1 that separation distance is a factor in UWB range error. This section uses the distance error model found in Section 3.1 to estimate the error and correct the measured range. The slope, m, of the linear separation distance models from each anchor were averaged and used for all three anchors during this error correction method. The y-intercept, b, was uniquely determined for each anchor by a calibration period at the beginning of each trial. The purpose of this was to adjust the model of the UWB to better fit the current scenario since it is known that errors change with many environmental factors.

A formula to determine the relationship between the measured range and the separation distance can be derived by substituting Equation 3.1 into Equation 3.2. The result gives

$$R = \frac{\hat{R} - b}{1 + m} \tag{4.1}$$

Equation 4.1 is used to correct each range measurement. The actual error can be computed by substituting Equation 4.1 into Equation 3.2, giving

$$E = \frac{m\hat{R} + b}{m + 1} \tag{4.2}$$

The process of correcting the range measurements first begins with the calibration period. While it is assumed that the slope, m, is known for all anchors (since all values were similar for each anchor-tag pair), the y-intercept, b, was assumed to be unknown (since it varies significantly between anchors). Therefore, the y-intercept was calculated using the average error during the calibration period such that

$$b = \bar{E}(1 + m) + m\bar{\hat{R}} \tag{4.3}$$

where $\bar{E}$ is the average error during the calibration period from the ground truth state propagation estimate and $\bar{\hat{R}}$ is the average range measurement during the calibration

period. After determining the value of b in the linear model for each anchor-tag pair, the model may be used to correct for the estimated error of each range measurement via Equation 4.1.

Ground truth measurements from the Vicon motion capture system were used to determine the *true* error of the uncorrected and corrected range measurements. The values in Table 4.1 denote the averaged mean and median absolute value of error for all three anchor-tag pairs. The columns labeled *percent change* present the change in mean and median absolute value of error from the original, uncorrected statistics. Negative values signify a decrease in error from the original measurements. Positive values denote an increase in error.

Table 4.1: Average mean and median range measurement error for each trial after error correction using the linear separation distance model.

Data Set	Trial	Mean (cm)	Mean (% change)	Median (cm)	Median (% change)
Front-Facing Antenna	Trial1	4.0	-77.6	3.5	-80.2
	Trial2	4.0	-77.4	3.6	-79.9
	Trial3	3.6	-80.1	3.1	-82.8
	Trial4	4.2	-76.8	3.6	-79.8
Right-Facing Antenna	Trial1	4.2	-73.4	3.7	-76.0
	Trial2	3.8	-75.2	3.3	-78.3
	Trial3	4.3	-72.1	3.7	-75.8
	Trial4	4.3	-71.4	3.7	-74.5

It can be seen that both mean and median error are greatly improved using this distance-based correction method. Maximum error is the result of a large spike in the individual range measurements present in all trials. It is a less important metric because it is hypothesized to be due to external factors — not of separation distance. Any change in the maximum error is somewhat arbitrary; therefore maximum error data have not been included in the analysis of any of the correction methods presented.

Figure 4-3 shows example plots from each data set of the individual error for each anchor. Figure 4-3a is a plot from trial 2 of the front-facing antenna data set and Figure 4-3b is from trial 1 of the right-facing antenna data set. These trials

were chosen as representative of all trials since the post-correction errors were overall similar to the averages. The large spike in error mentioned in the previous paragraph can be seen in Figures 4-3a and 4-3b between 80 to 100 seconds for various UWB anchors.

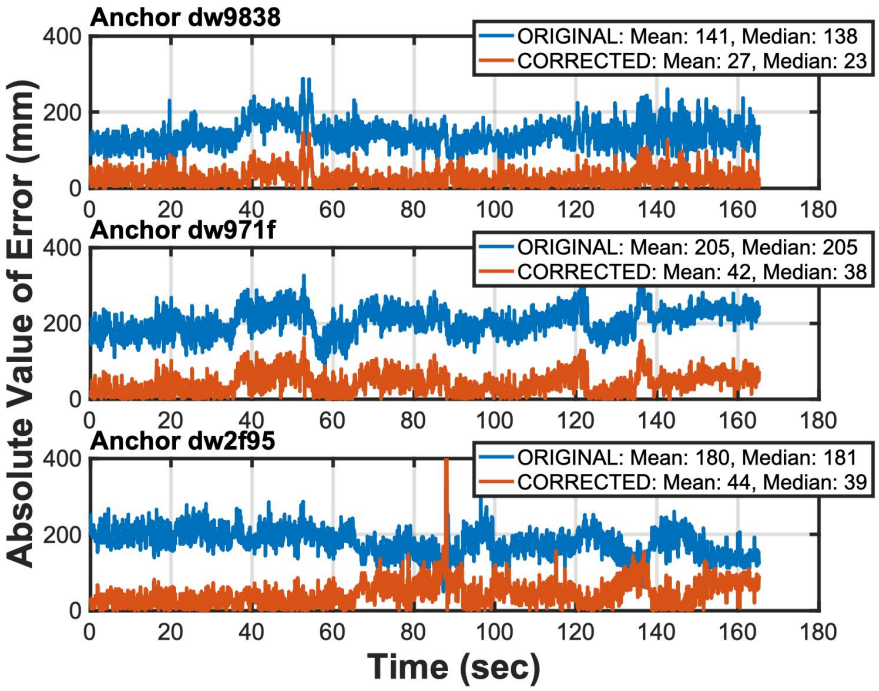

(a) Trial 2 of the front-facing antenna data set

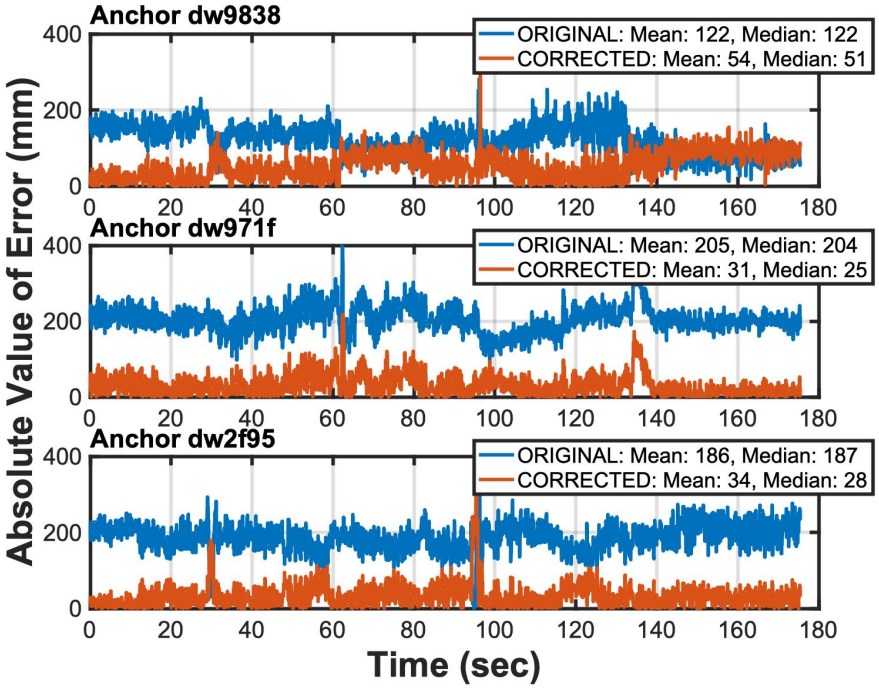

(b) Trial 1 of the right-facing antenna data set

Figure 4-3: Error of the range measurements from each anchor-tag pair before and after error correction using the linear separation distance model.

4.3 Error Correction using Relative Angle

While Section 4.2 addressed correcting error according to the estimated separation distance between the anchor and tag, Section 4.3 uses the relative angle elliptical models developed in Section 3.2 to estimate and correct for measurement error. One major assumption in this method is that both the size and shape of the angular elliptical model are consistent across all separation distances for robot localization. This is a known simplification since Section 3.1 showed that at the very minimum, the error between two modules with a relative angle of 180 degrees is dependent upon separation distance.

The elliptical models for each anchor-tag pair from Section 3.2 were used. To create a more generalized model, the parameters describing each of the three elliptical models were individually averaged to form a single model. Since it is known that error changes as a function of environmental factors, the calibration period is used to scale the elliptical models to better fit the current trial. During the calibration period, the robot's state — as determined by the propagated equations of motion — is considered to be ground truth, as it was in the distance-based error correction method. The relative angle between the anchor and tag is calculated by Equation 3.3.

For the entire trial, the heading of the anchor's antenna, α, is known, and the heading of the tag, β, is calculated by the propagated equations of motion from Section 2.1.[1] During the calibration period, to sort the calculated percent error as a function of relative angle, a 360 degree circle was sliced into bins of 0.05 radians, or roughly 2.9 degrees, each. For each anchor, the estimated percent range error was sorted into the correct bin. At the end of the calibration period, the average error for each bin was computed. This was done to identify the relationship between relative angle and error for the current trial in order to properly scale the elliptical models.

From the calibration period, a scaling factor was determined to adjust the size of the elliptical relative angle versus percent error models for the current conditions. The

[1]Potential higher-accuracy methods of estimating the robot's heading during localization are discussed in 5.2.

scaling factor, S, is essentially equivalent to the percent error during the calibration period divided by the expected percent error during the calibration period, according to the generalized ellipse model. More specifically, the formula for the scaling factor is

$$S = \frac{\bar{\epsilon}}{\tilde{\epsilon}_g} \tag{4.4}$$

where $\bar{\epsilon}$ is the mean percent error for all populated bins of the calibration period and $\tilde{\epsilon}_g$ is the expected percent error according to the generalized ellipse model using the median relative angle during the robot's rotation. Finally, for each range measurement, the estimated percent error, ϵ, is determined by finding the predicted error from the elliptical model given the relative angle and multiplying it by the scaling factor such that

$$\epsilon = S\epsilon_a \tag{4.5}$$

where ϵ_a is the predicted error from the unscaled, generalized ellipse model for the current relative angle. Percent error is equivalent to the range error, E, divided by the separation distance, R, such that

$$\epsilon = \frac{E}{R} \tag{4.6}$$

The value of ϵ is estimated from Equation 4.5. By combining Equation 3.1 with Equation 4.6, an estimate for the separation distance, R, can be calculated as

$$R = \frac{\hat{R}}{1 + \epsilon} \tag{4.7}$$

The percent error versus relative angle data from the calibration period are shown in Figure 4-4, alongside the calibrated elliptical models. Figure 4-4a shows the data from Trial 2 of the front-facing antenna data set, and Figure 4-4b shows the data from Trial 1 of the right-facing antenna data set. Although the measured relative angles of

each anchor-tag pair are different for each data set, the general sizes of the angular models are very similar for each anchor between data sets. For anchor *dw971f*, the lowest error is seen at a relative angle of zero degrees for Figures 4-4a and 4-4b. This is the same trend as seen in the static angular data in Section 3.2, which suggests that adjusting the current elliptical model to include the lower error at and around zero degrees may further decrease range error and improve localization accuracy.

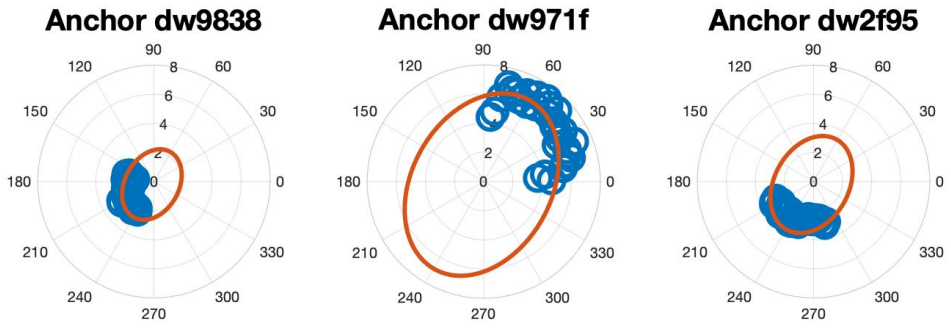

(a) Trial 2 of the front-facing antenna data set

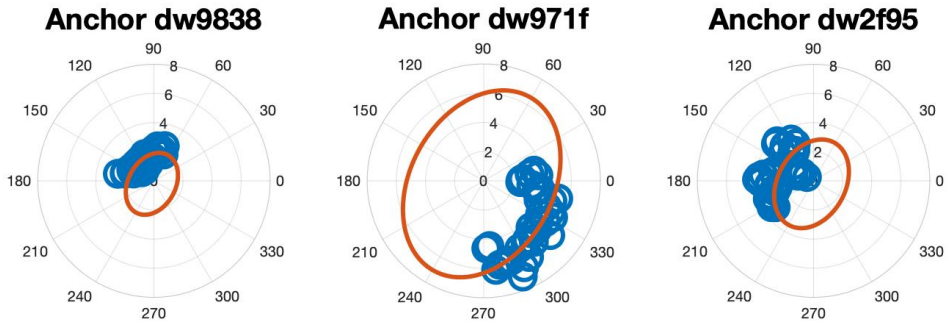

(b) Trial 1 of the right-facing antenna data set

Figure 4-4: Using the data from the calibration period, the relationship between relative angle and percent error is plotted in blue. The calibrated elliptical models are overlaid on the raw data.

Ground truth measurements from the Vicon motion capture system were used to determine the error of the initial measurements and the error of the angle-corrected measurements. Table 4.2 presents the mean and median range error for each trial after angle-based error correction. Comparing these data to the distance-corrected results from Table 4.1, the error is lower for the distance-based correction. Nonetheless, the

angle-based correction decreases the mean error by an average of 65.2% for all trials of both data sets.

Table 4.2: Average mean and median range measurement error for each trial after error correction using the elliptical relative angle model.

Data Set	Trial	Mean (cm)	Mean (% change)	Median (cm)	Median (% change)
Front-Facing Antenna	Trial1	6.9	-63.5	6.2	-66.9
	Trial2	6.7	-64.5	5.6	-70.4
	Trial3	6.4	-66.0	5.5	-71.1
	Trial4	6.3	-66.1	5.6	-69.9
Right-Facing Antenna	Trial1	6.4	-63.8	5.7	-68.2
	Trial2	7.0	-60.9	6.5	-64.1
	Trial3	8.5	-52.9	8.4	-55.1
	Trial4	6.7	-62.2	6.3	-65.3

A more detailed understanding of the error for each anchor-tag pair can be seen in Figure 4-5, which shows the individual error for all anchors of trial 2 of the front-facing antenna data set and trial 1 of the right-facing antenna data set. Again, these trials were chosen as representative of all trials. This figure is comparable to Figure 4-3, which shows the individual improvements in range error for distance-based correction. From the size of the elliptical model, it can be seen that anchor *dw971f* has comparatively greater predicted error during the calibration period. This is opposite of what we would expect since the distance from the tag to anchor *dw971f* is shorter than the distance to either of the other anchors. This could be partially due to the accuracy of the robot's equations of motion, as it was seen in Section 2.1 the majority of the error was along the same direction as the the line between anchor *dw971f* and the tag.

This correction method does not improve anchor *dw971f*'s range error as much as either of the other two anchors. The reason for this is unclear, though it could be due to the relative angle during the calibration period since neither of the other two anchors have similar relative angles in the calibration period of either data set. It could also be connected to the separation distance since during the calibration period

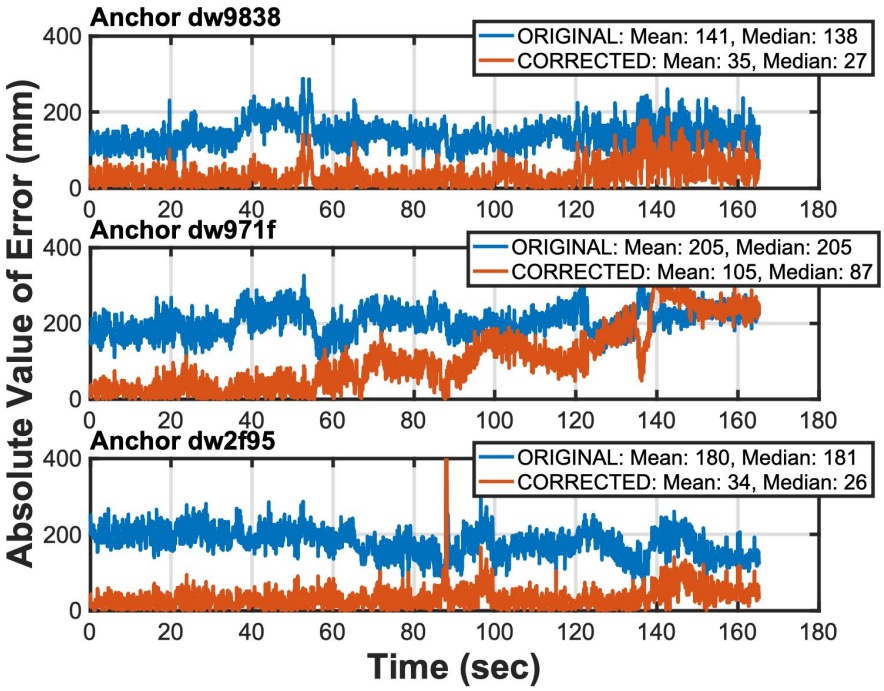

(a) Trial 2 of the front-facing antenna data set

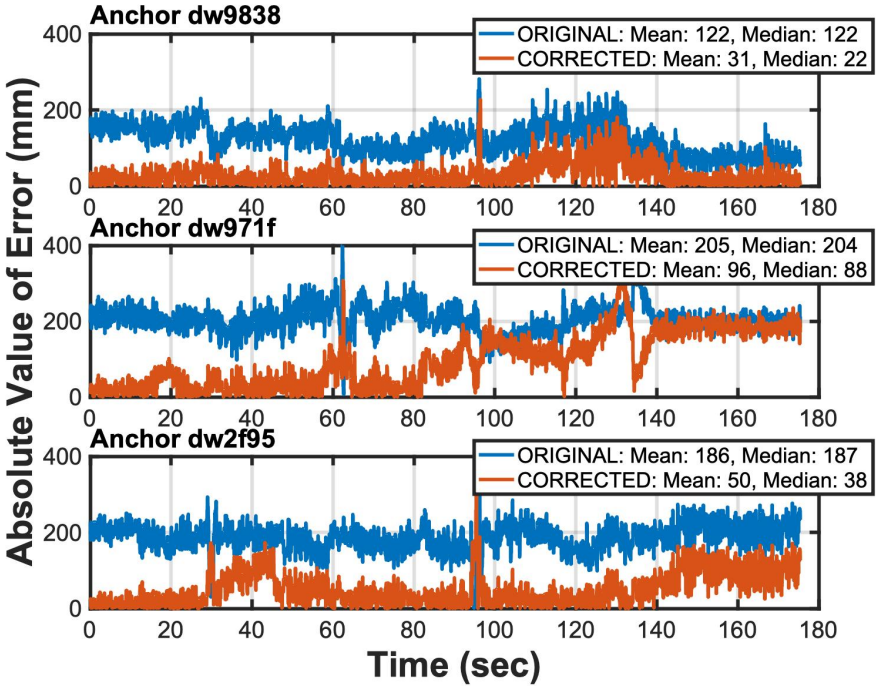

(b) Trial 1 of the right-facing antenna data set

Figure 4-5: Error of the range measurements from each anchor-tag pair before and after error correction using the elliptical relative angle model.

anchor *dw971f* has a separation distance equivalent to about half the separation distance of the other two anchors.

4.4 A Fused Model for Correcting Error

Sections 4.3 and 4.3 used separation distance and relative angles as separate methods to correct for error of UWB range measurements. In this section, these two error models are fused to provide a more thorough description of error which can correct range measurements more consistently and produce localization results with lower error than either of the former models. The key assumption underlying this fused model is that the elliptical *shape* of the angular error model is constant for all separation distances, though the *size* of the ellipse scales linearly with the expected percent error from the calibrated separation distance model.

The parameters describing the linear distance error model and the elliptical angular error model were found using the same methods as described in Section 4.2 and Section 4.3, respectively. The way that these two models were tied together was by the distance error ratio, r. The purpose of the distance error ratio is to constantly adjust the size of the ellipse model to reflect the current estimated separation distance. Recall from Section 4.3 that the scaling factor essentially scales the generalized ellipse model to fit the data during the calibration period. Therefore, the distance error ratio is a fraction of the expected percent error of the current measurement divided by the expected error of the calibration period, both according to the linear separation distance model. The calculation for r is

$$r = \frac{\epsilon_d}{\bar{\epsilon}_d} \tag{4.8}$$

where ϵ_d and $\bar{\epsilon}_d$ are the percent errors predicted by the calibrated linear distance model. The formula for ϵ_d is calculated by dividing Equation 4.2 by Equation 4.1,

which results in

$$\epsilon_d = \frac{m\hat{R} + b}{\hat{R} - b} \tag{4.9}$$

While ϵ_d is calculated directly via Equation 4.9, $\bar{\epsilon}_d$ is calculated via Equation 4.9 by setting $\hat{R}$ equal to the average range measurement of the anchor during the calibration period such that

$$\bar{\epsilon_d} = \frac{m\bar{\hat{R}} + b}{\bar{\hat{R}} - b} \tag{4.10}$$

where $\bar{\hat{R}}$ is the average range measurement of the anchor during the calibration period.

Then, the predicted percent error from the fused model can be calculated by

$$\epsilon = r\epsilon_a \tag{4.11}$$

where ϵ is the final predicted percent error according to the fused model and ϵ_a is the predicted percent error according to the angular model. Finally, the corrected range is calculated by Equation 4.7.

It was found that this method of error correction resulted in lower average mean and median range error for all trials when compared to the distance- and angle-based error correction methods. The mean and median measurement error for each trial using the fused model to correct error can be seen in Table 4.3.

The fused error correction method was more consistent in its corrections than either of the previous methods, though it did not always produce the lowest error for each individual anchor-tag pair. Plots of the individual errors from the selected, representative trials of both data sets showing the results of the fused model error correction are presented in Figure 4-6.

Table 4.3: Average mean and median range measurement error for each trial after error correction using the fused distance-angle model.

Data Set	Trial	Mean (cm)	Mean (% change)	Median (cm)	Median (% change)
Front-Facing Antenna	Trial1	2.9	-83.2	2.5	-85.4
	Trial2	2.9	-83.0	2.4	-85.9
	Trial3	2.7	-84.5	2.2	-87.3
	Trial4	3.0	-83.0	2.4	-85.9
Right-Facing Antenna	Trial1	3.2	-80.7	2.8	-83.1
	Trial2	3.4	-78.7	2.9	-81.4
	Trial3	3.9	-76.4	3.2	-80.2
	Trial4	3.2	-79.7	2.7	-83.1

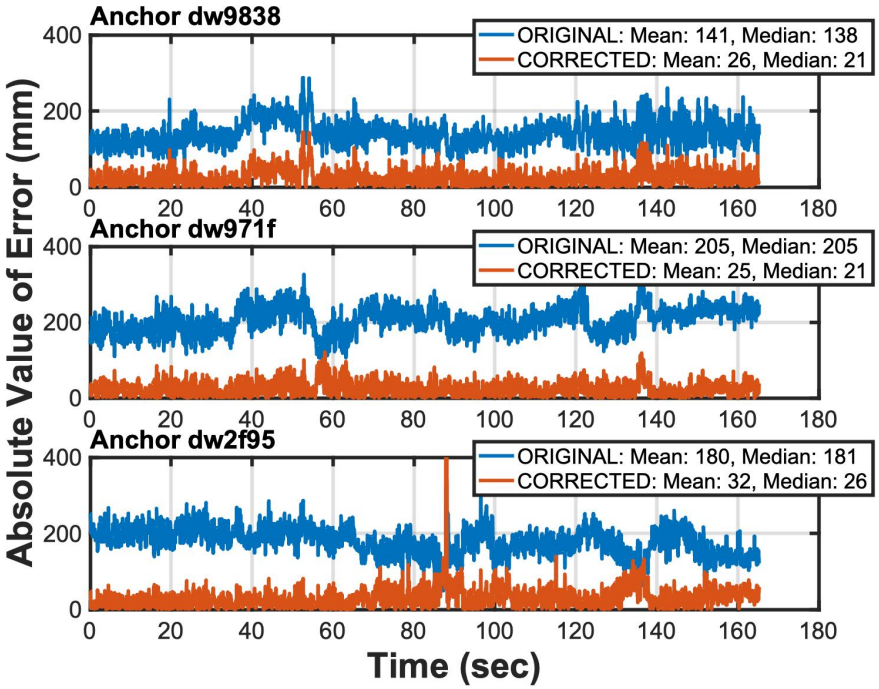

(a) Trial 2 of the front-facing antenna data set

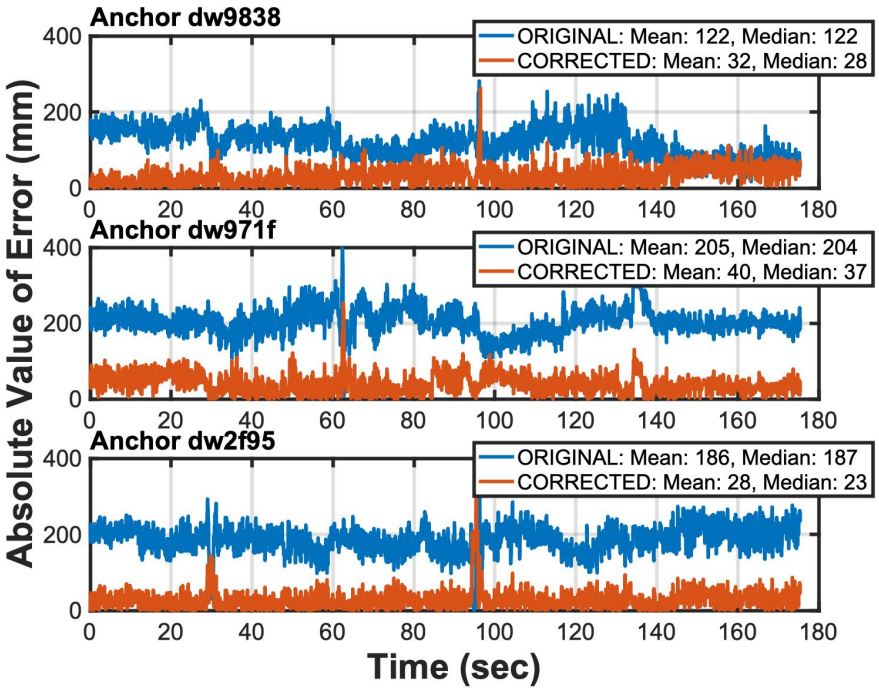

(b) Trial 1 of the right-facing antenna data set

Figure 4-6: Error of the range measurements from each anchor-tag pair before and after error correction using the fused distance-angle model.

4.5 Comparing Correction Methods

All three error correction methods substantially decreased the individual range errors, however, by directly comparing the range and localization error from each data set, statistics show that the fused model is superior at decreasing localization error. Table 4.4 presents the average mean and median range error values for all anchors in each data set. Overall, the error was highest for the angle-corrected measurements, and lowest for the fused model-corrected measurements.

Table 4.4: Average error mean and median error of range measurements for all trials of each data set and correction method. Bold values indicate the correction method resulting in the lowest error or the greatest decrease in error for each data set.

		Mean		Median	
Data Set	Method	(cm)	(% change)	(cm)	(% change)
Front-Facing Antenna	None	17.8	–	17.6	–
	Distance	3.9	-78.0	3.5	-80.7
	Angle	6.6	-65.0	5.7	-69.6
	Fused	**2.8**	**-83.4**	**2.4**	**-86.1**
Right-Facing Antenna	None	16.8	–	16.8	–
	Distance	4.1	-73.0	3.6	-76.2
	Angle	7.1	-60.0	6.7	-6.2
	Fused	**3.4**	**-78.9**	**2.9**	**-82.0**

Figure 4-7 presents the error for individual anchors of trial 2 of the front-facing antenna data set for all three error correction methods alongside each other.[2] To understand the qualitative patterns of the error correction method, first consider Figure 4-7b, showing the angle-based correction method. It can be seen that the error increases toward the end of the robot's path because the UWB error model is less accurate since the distance between the anchor and tag has changed. This is especially noticeable for anchor *dw971f*. Next, looking at Figure 4-7a, it can be seen that there is no longer an increasing linear trend in the corrected error. That trend is corrected when adjusting the predicted error based on separation distance. Finally,

[2]The three plots in Figure 4-7 are copies of the plots in Figures 4-3a, 4-5a, and 4-6a.

consider Figure 4-7c which uses a fused model to correct for error due to distance and angle simultaneously. Using the fused model, the corrected error magnitude is more consistent across the robot's entire trajectory and the sections of higher absolute value error in Figure 4-7a are brought closer to zero.

Looking at the decrease in error of specific anchors, the errors of anchors *dw2f95* and *dw9838* are only about 1 cm different for the angle- and distance-based correction methods. While anchor *dw2f95*'s error is lower using the angle-based error correction, anchor *dw9838*'s error is lower when correcting via distance. The error of anchor *dw971f* after angle-based correction is 2.5 times higher than when using distance-based correction. From the distance-based to the fused model corrections, it can be seen that the biggest improvement is seen for anchor *dw971f*. This is hypothesized to be because anchor *dw971f* is the only anchor in this trial whose relative angles are not close to the 180 degrees which the distance-based correction assumes.

While the errors after distance-based correction are close to the errors after the fused model correction, this is not the case for localization error. Least squares estimation (detailed in Section 2.2) was used to localize the tag with (1) the uncorrected range measurements, (2) the distance-corrected measurements, (3) the angle-corrected measurements, and (4) the fused model-corrected measurements. The mean and median localization error is shown for each data set in Table 4.5.

It can be seen that although both the distance- and angle-based correction methods greatly decrease individual range measurement error, the overall localization error does not always decrease. The angle-based correction method in particular causes an *increase* in localization error for all trials. The distance-based correction does slightly better, with mean values decreasing in five of eight trials and increasing in only three of them. The only error correction method presented in this paper which decreases the range-based least squares localization error for every trial is the fused model correction. Furthermore, even in the trials for which the distance-based correction decreases localization error, the fused model's localization error is equal or lower. In fact, only considering the trials in which the distance-based correction decreases localization error, the fused-model localization error is still nearly 20% lower on average.

Table 4.5: Average mean and median localization error using least squares estimation for all trials of each data set and correction method. Bold values indicate the correction method resulting in the lowest error or the greatest decrease in error for each data set.

Data Set	Method	Mean (cm)	Mean (% change)	Median (cm)	Median (% change)
Front-Facing Antenna	None	9.9	–	8.8	–
	Distance	10.3	+3.8	8.8	0.0
	Angle	11.9	+20.2	10.5	+18.8
	Fused	**6.7**	**-30.2**	**5.3**	**-39.5**
Right-Facing Antenna	None	12.1	–	11.6	–
	Distance	8.8	-27.8	7.6	-34.6
	Angle	14.1	+16.8	13.2	+14.4
	Fused	**7.3**	**-39.7**	**6.3**	**-45.9**

The change in localization error can be seen in Figure 4-8, which shows plots of the uncorrected localization and the localization after all three correction methods for trial 2 of the front-facing antenna data set. Although the mean and median range errors are greatly reduced by all three correction methods (Table 4.5), the mean localization error for this trial actually increases for the distance- and angle-based corrections, by 4.1 and 25.8% respectively. On the other hand, the fused model *decreases* the mean localization error by 25.8%.

Looking qualitatively at the plots, the uncorrected localization, Figure 4-8a, has significant error in the X direction when the robot moves parallel to the y-axis. In the angle-corrected localization plot, Figure 4-8c, it is interesting to see the increase in localization error as the robot gets toward the end of its path. This is likely a result of the increase in error of the corrected measurements as the robot moves away from its starting point and the elliptical error model becomes less accurate. Visually, the distance-based correction, Figure 4-8b, seems to do much better than both the uncorrected and angle-corrected localization, however, this is not the case. While the distance-corrected localization does have lower error than the angle-corrected localization, it has higher error than the uncorrected measurements. The fused model, Figure 4-8d, has much lower error than the three other methods. For example, one

place this can be seen is by the strong correlation between the localization estimate and ground truth position while the robot is traveling along the x-axis.

To see why overall lower range error may not always translate to lower localization error, consider the example given in Figure 4-9. Figure 4-9a depicts a scenario where all three range measurements have high error, but their errors seem to "cancel" to result in a localization estimate with low error. However, Figure 4-9b shows a scenario where two of the measurements are very accurate and only one has high error, which results in a localization estimate with low accuracy. In particular, the pattern of higher error leading to lower accuracy can be visually observed when using the angle-based correction method. By comparing Figure 4-7b and 4-8c, it can be seen that even though the mean range error decreased with respect to the original measurements, the localization error did not. In this specific case, anchor *dw971f* is disproportionately responsible for the measurement error, in a scenario similar to the bottom left range measurement in Figure 4-9b. The localization error summary in Table 4.5 and the theoretical example given in Figure 4-9 is evidence that it is not enough to decrease the mean range errors, but decreasing range errors in a *consistent* manner is vital to improving localization error.

After performing least squares estimation, the estimated positions were used in a Kalman Filter to smooth out the position estimate. The linear state space model is described by Equation 2.1 and the measurement equation is described by Equation 2.4. The covariance matrices for the process and measurement noise were kept constant throughout all trials. The average mean and median localization error of each data set are shown in Table 4.6. All percent changes were computed with respect to the uncorrelated localization error using the Kalman Filter.

The same pattern as previously seen occurs when using a Kalman filter for localization as well: decreasing measurement error does not always decrease localization error. As was seen in the least squares localization estimate, the fused model was the only model to decrease mean error in all trials of both data sets. A plot of the error over time for a representative trial, trial 2 of the front-facing antenna data set, can be seen in Figure 4-10.

Table 4.6: Average mean and median filtered localization error for all trials of each data set and correction method. Bold values indicate the correction method resulting in the lowest error or the greatest decrease in error for each data set.

Data Set	Method	Mean (cm)	Mean (% change)	Median (cm)	Median (% change)
Front-Facing Antenna	None	8.6	–	9.2	–
	Distance	9.1	+9.0	8.4	-8.2
	Angle	9.7	+13.5	8.7	-3.7
	Fused	**4.2**	**-51.1**	**3.4**	**-62.7**
Right-Facing Antenna	None	11.6	–	12.0	–
	Distance	6.3	-45.5	6.1	-48.1
	Angle	11.7	+2.2	11.5	-3.7
	Fused	**4.6**	**-60.9**	**4.0**	**-66.9**

In Figure 4-10, it is shown that the localization error of the angle-corrected ranges increases with time. The distance-corrected and fused model-corrected error follow a similar trend, although the distance-corrected localization error is greater for almost every time step. For each error correction method, the resulting localization estimates for this representative trial are plotted in Figure 4-11.

Overall, it can be seen that the localization error is further improved with a Kalman Filter as compared with the raw least squares estimate. For the fused model specifically, the mean error of all trials is decreased by an additional 2.5 cm, or approximately an additional 35%.

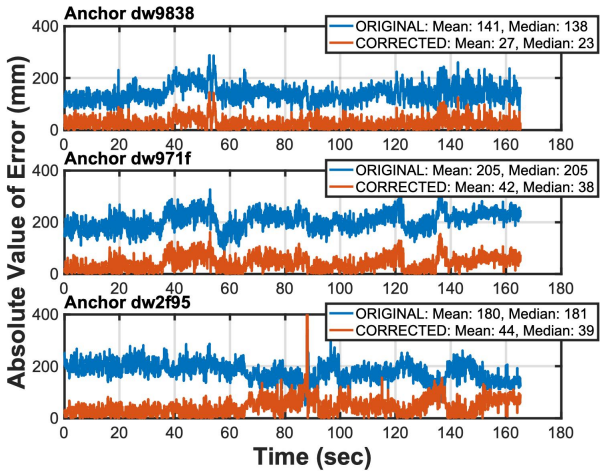

(a) Distance-based correction

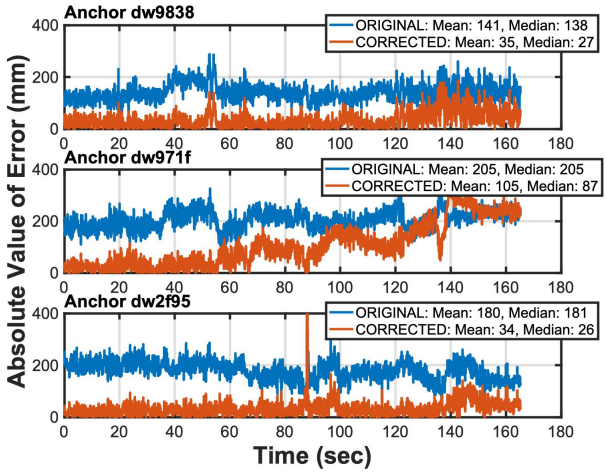

(b) Angle-based correction

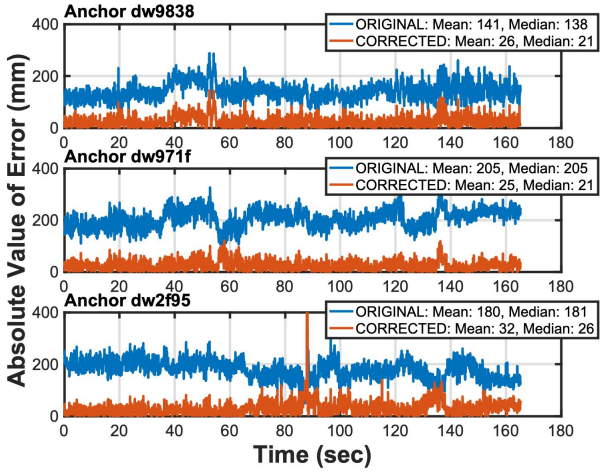

(c) Fused model correction

Figure 4-7: Error of the measurements from each anchor-tag pair of Trial 2 of the front-facing antenna data set for all three error correction methods.

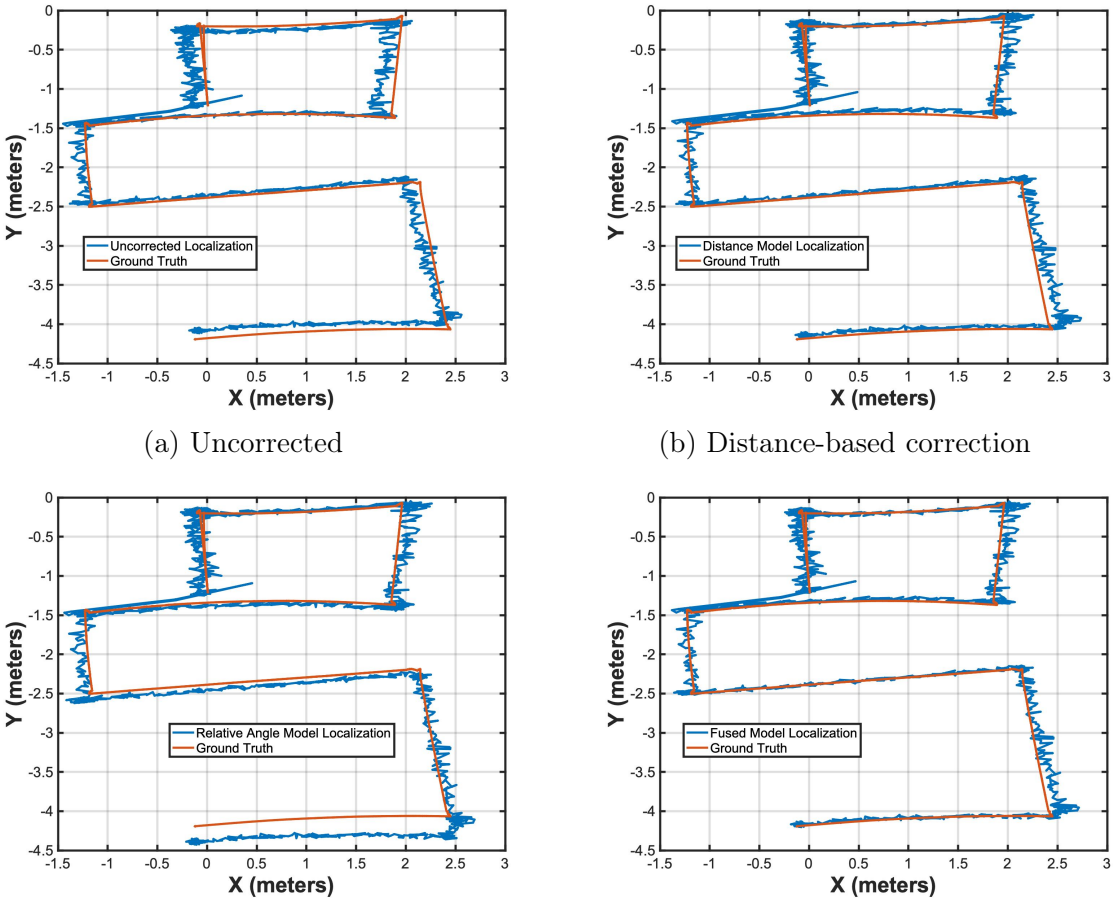

(a) Uncorrected

(b) Distance-based correction

(c) Angle-based correction

(d) Fused model correction

Figure 4-8: The resulting localization using least squares estimation for Trial 2 of the front-facing antenna data set for the uncorrected measurements and for all three error correction methods.

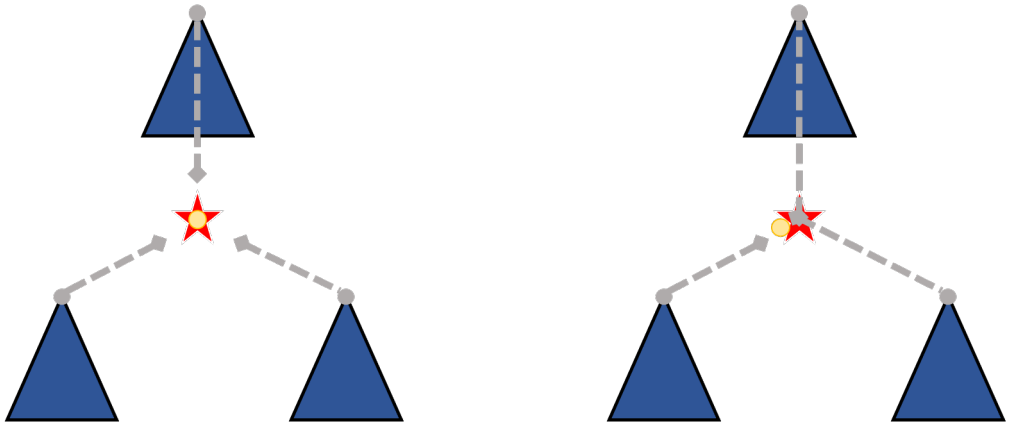

(a) A set of equally bad range measurements could lead to good localization.

(b) One poor range measurement can result in poor localization.

Figure 4-9: Range measurements with high error can result in good localization and measurements with comparatively lower error can result in poor localization. The red star is the actual position of the tag, the grey lines represent the range measurements, and the yellow circle is the theoretical localization estimate given those range measurements.

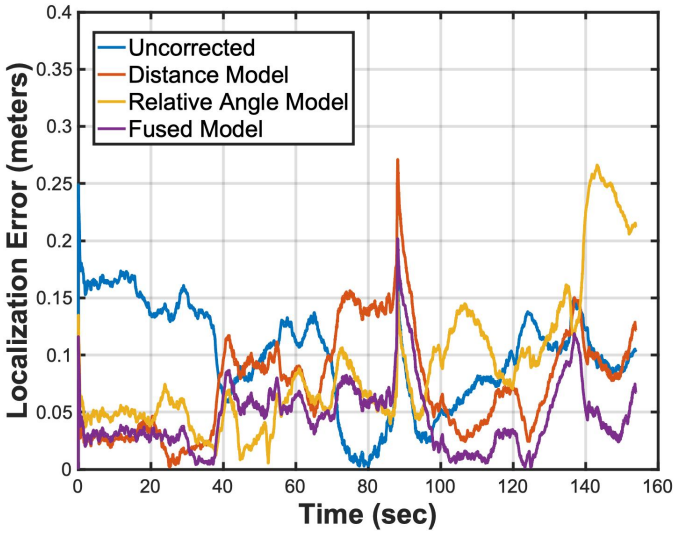

Figure 4-10: The localization error over time for Trial 2 of the front-facing antenna data set for all correction methods, smoothed with a Kalman Filter.

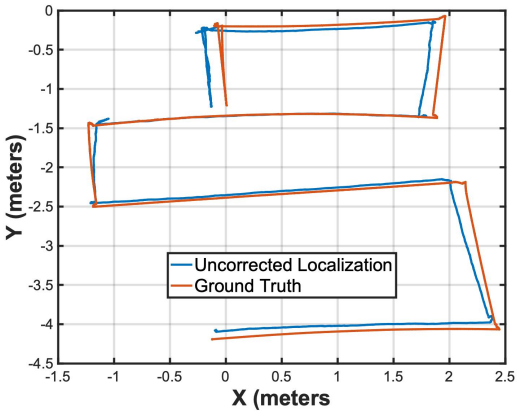

(a) Uncorrected

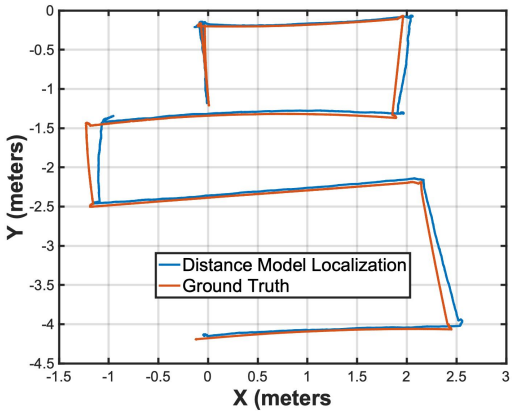

(b) Distance-based correction

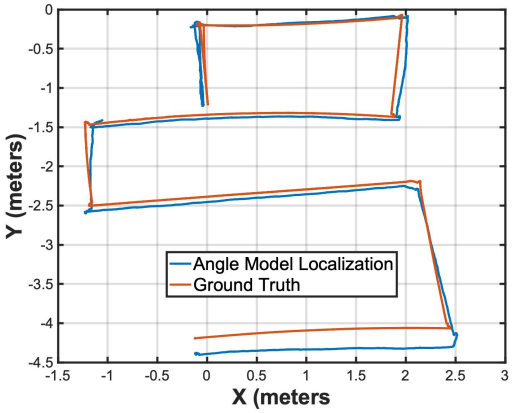

(c) Angle-based correction

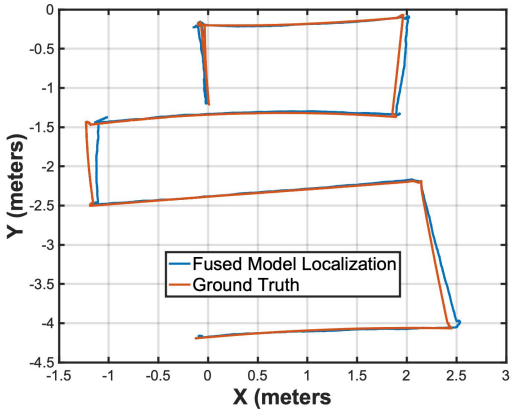

(d) Fused model correction

Figure 4-11: The resulting localization, smoothed with a Kalman Filter, for Trial 2 of the front-facing antenna data set for the uncorrected measurements and for all three error correction methods.

Chapter 5

Conclusion

5.1 Summary of Contribution

With the goal of improving accuracy of range-based localization using UWB, the data presented builds upon the current understanding of UWB error, developing simple models to describe measurement error as a function of separation distance and relative angle between an anchor and tag. This work applies these models to robot localization to predict and correct for measurement error, resulting in approximately 80% lower range error, 35% lower localization error when using a least squares approach, and 56% lower localization error when smoothing the trajectory with a Kalman Filter.

Two sets of data with static anchors and static tags were taken separately to explore the relationship between separation distance and relative angle with UWB range error. It was seen that the magnitude of range error increased with separation distance in a generally linear trend. Using least squares estimation, a line was fit to the data, and the resulting slopes of the lines were consistent for all three anchor-tag pairs. When plotting percent error versus relative angle in polar coordinates, the error was shown to be elliptical in shape, therefore an ellipse was fit to the data. It was found that the eccentricities and ellipses of all three anchor-tag pairs were similar.

The three different methods of error correction developed and analyzed in this work include: (1) a method based on estimating error from only separation distance, (2) a method based on estimating error from only relative angle, and (3) a method

which fuses the separation distance and relative angle models together to estimate error (the *fused model*). All three methods start with distance and/or relative angle error models generalized from the stationary data models. Since UWB range error is affected by many uncontrollable environmental factors, all three methods also use a short calibration period before localization starts to adjust the models to fit the error patterns of the current environment.

Two sets of range data of a robot driving a lawnmower pattern were used to validate the three different correction methods. The only difference between the two data sets was the orientation in which the UWB tag was mounted on the robot. All three correction methods significantly lowered the error of the individual range measurements. Averaged over all trials of both data sets, the distance-based correction decreased the mean error by 75.5%, the angle-based correction by 62.5%, and the fused model by 81.6%. In general, all correction methods resulted in greater decreases in individual range errors for the front-facing antenna data set than for the right-facing antenna data set. The only difference in the two data sets is the orientation of the tag on the robot, suggesting that perhaps there are more intricacies to the relationship between relative angle and range error than were discovered in this work.

To test localization accuracy, the corrected measurements from all three methods were used to localize the robot using least squares estimation. Although individual range errors were significantly decreased using all three correction methods, this was not so with localization error. The distance-based correction method only decreased the mean localization error in five of the eight trials, and the angle-based correction method decreased the mean localization error for none of the trials. The fused model outperformed the other two, with 25.8 to 44.4% decreases in error for all trials. These results show that decreases in individual range error do not always lead to higher accuracy localization.

Next, the range based localization estimate was smoothed using a Kalman Filter. Using the Kalman filter, the distance-based error correction only improved the mean localization error for half the trials, and the angle-based error correction improved three of eight. The fused model error correction significantly improved the localization

accuracy of all trials, lowering the error by over 50% in six of the eight.

Range-based localization is a well-researched area and frequently used for everyday tasks such as tracking and navigation. While there are many technologies available to estimate distance, many require a direct line-of-sight. UWB is a contemporary technology, lauded for its ability to transmit a signal through solid objects to measure distance between antennas. Although this quality makes UWB a favorable option for indoor and cluttered environments, it does come with challenges. Most notably, the error is sensitive to many environmental factors and even to the exact geometry between antennas. This work focuses on modeling and correcting for UWB error related to separation distance and the relative angle between antennas, contributing to the understanding of UWB error patterns.

5.2 Future Research

This thesis confirmed the distance and angle dependence of UWB range measurements with the Decawave DWM1001 module and developed a fused error model to correct for this error. A possible next step in this work is to extrapolate these findings to new scenarios. For example, since these generalized models were developed with data from the same UWB modules as were used for testing, the robustness of these models could be tested with robot experiments in new environments, under NLOS conditions, or even with other brands of UWB antennas. Data taken from additional DWM1001C modules give us confidence that the models could be extrapolated to other Decawave modules.

These data focused on separation distances above 3000 mm. However, despite potential near-field antenna effects, it is believed that the three error correction methods developed would decrease the range error of measurements taken at distances shorter than 3000 mm as well. The only anticipated constraint is that the error correction methods would result in a greater reduction of error if the robot calibration period began with a separation distance of 3000 mm or greater from all anchors. Future research could be done to confirm this hypothesis.

Additionally, our model and experiments are based on motion of a robot in the plane of the UWB modules. Further research and analysis would be necessary for applications that would involve three-dimensional motion, such as experiments with unmanned aerial vehicles.

A potential application of this work could be for error correction of cooperative localization. Sometimes localization capabilities are required in situations where access to the area to pre-place stationary UWB anchors is not possible due to safety (natural disasters, collapsed buildings, military applications) or accessibility (lava tubes, outer space, underwater). Cooperative localization uses UWB modules on multiple robots for short-distanced relative localization of the individual robots within the swarm. The error correction methods developed in this thesis could be applied to cooperative localization of robot swarms to correct for range measurement error.

Another area of future work could consider the estimate of the robot's heading during localization. In this thesis, the propagated equations were used to estimate the heading of the tag at all time steps along the trajectory. The model validation in Section 2.1 showed the mean absolute value of heading error to be 1.6 degrees, however this accuracy of measurement is not guaranteed for all robots or environments. Future work could benefit from improving the accuracy of the heading estimate. One potential way to do this is to mount a second tag on the robot so that an estimate for the heading may be calculated by range measurements. This estimate could be combined with the heading estimate from the robot model using a Kalman Filter (or similar) to improve heading accuracy. A more complex solution may include the development of an algorithm which combines each range measurement with the relative angle error models to find a probability distribution which describes the robot's heading at each time step, solving both the error prediction and heading estimation problems simultaneously.

Bibliography

[1] K. Bregar, T. Javornik, A. Hrovat, M. Mohorčič, and G. Kandus. Passive ultra-wideband coarse localization and activity detection system for assisted living. In *2019 23rd International Conference on Applied Electromagnetics and Communications (ICECOM)*, pages 1–5, 2019.

[2] K. Bregar and M. Mohorčič. Improving indoor localization using convolutional neural networks on computationally restricted devices. *IEEE Access*, 6:17429–17441, 2018.

[3] A. Cazzorla, G. De Angelis, A. Moschitta, M. Dionigi, F. Alimenti, and P. Carbone. A 5.6-ghz UWB position measurement system. *IEEE Transactions on Instrumentation and Measurement*, 62(3):675–683, 2013.

[4] L. Clark, C. Andre, J. Galante, B. Krishnamachari, and K. Psounis. TEAM: Trilateration for exploration and mapping with robotic networks, 2020.

[5] G. De Angelis, A. Moschitta, and P. Carbone. Positioning techniques in indoor environments based on stochastic modeling of UWB round-trip-time measurements. *IEEE Transactions on Intelligent Transportation Systems*, 17(8):2272–2281, 2016.

[6] Decawave. *APS013 APPLICATION NOTE: The implementation of two-way ranging with the DW1000*, 2015.

[7] Decawave. *DW1000 Datasheet*, 2016.

[8] Decawave. *DWM1001 System Overview And Performance*, 2018.

[9] J. Fu, Y. Fu, and D. Xu. Application of an adaptive UKF in UWB indoor positioning. In *2019 Chinese Automation Congress (CAC)*, pages 544–549, 2019.

[10] P. Furgale, T. Barfoot, and G. Sibley. Continuous-time batch estimation using temporal basis functions. *The International Journal of Robotics Research*, 34:2088–2095, 05 2012.

[11] P. Furgale, C.H. Tong, T.D. Barfoot, and G. Sibley. Continuous-time batch trajectory estimation using temporal basis functions. *International Journal of Robotics Research*, 34(14):1688–1710, 2015.

[12] O. Gal. MATLAB central file exchange, fit_ellipse, 2021.

[13] J. González-Jiménez, J. Blanco, C. Galindo, A. Ortiz-de Galisteo, J.A. Fernández-Madrigal, F. Moreno, and J. Martínez. Mobile robot localization based on ultra-wide-band ranging: A particle filter approach. *Robotics and Autonomous Systems*, 57:496–507, 05 2009.

[14] M. Hamalainen, P. Pirinen, J. Iinatti, and A. Taparugssanagorn. UWB supporting medical ICT applications. In *2008 IEEE International Conference on Ultra-Wideband*, volume 3, pages 15–16, 2008.

[15] J. Keiller. Localization for teams of autonomous ground vehicles. Master's thesis, University of Illinois at Urbana-Champaign, 2019.

[16] J. Khodjaev, Y. Park, and A. Malik. Survey of NLOS identification and error mitigation problems in UWB-based positioning algorithms for dense environments. *Annales des Télécommunications*, 65:301–311, 06 2010.

[17] D. Lamensdorf and L. Susman. Baseband-pulse-antenna techniques. *IEEE Antennas and Propagation Magazine*, 36(1):20–30, 1994.

[18] A. Ledergerber and R. D'Andrea. Ultra-wideband range measurement model with Gaussian processes. *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 1929–1934, 2017.

[19] S. Maranó, W. M. Gifford, H. Wymeersch, and M. Z. Win. NLOS identification and mitigation for localization based on UWB experimental data. *IEEE Journal on Selected Areas in Communications*, 28(7):1026–1035, 2010.

[20] E. Pancera, T. Zwick, and W. Wiesbeck. Spherical fidelity patterns of UWB antennas. *IEEE transactions on antennas and propagation*, 59(6 PART 2):2111–2119, 06 2011.

[21] A. Prorok and A. Martinoli. Accurate indoor localization with ultra-wideband using spatial models and collaboration. *The International Journal of Robotics Research*, 33(4):547–568, 2014.

[22] S. Rutherford. What is UWB and why are Apple and Samsung suddenly so interested in it? *Gizmodo*, 12 2020.

[23] M. Sharma, C. Parini, and A. Alomainy. Influence of antenna alignment and line-of-sight obstruction on the accuracy of range estimates between a pair of miniature UWB antennas. *2015 9th European Conference on Antennas and Propagation (EuCAP)*, pages 1–5, 2015.

[24] J. Sidorenko, V. Schatz, N. Scherer-Negenborn, M. Arens, and U. Hugentobler. Decawave UWB clock drift correction and power self-calibration. *Sensors*, 19, 07 2019.

[25] P. Tabaghi, I. Dokmanić, and M. Vetterli. On the move: Localization with kinetic euclidean distance matrices. *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4893–4897, 2019.

[26] P. Tabaghi, I. Dokmanić, and M. Vetterli. Kinetic euclidean distance matrices. *IEEE Transactions on Signal Processing*, 68:452–465, 2020.

[27] J. Tiemann, J. Pillmann, and C. Wietfeld. Ultra-wideband antenna-induced error prediction using deep learning on channel response data. In *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*, pages 1–5, 2017.

[28] Steven Tully, G. Kantor, and H. Choset. Leap-frog path design for multi-robot cooperative localization. In *FSR*, 01 2009.

[29] D. Veit, M. Gadringer, and E. Leitgeb. Impact of UWB antennas on ranging accuracy. In *2020 14th European Conference on Antennas and Propagation (EuCAP)*, pages 1–5, 2020.

[30] Z. Xiao, H. Wen, A. Markham, N. Trigoni, P. Blunsom, and J. Frolik. Non-line-of-sight identification and mitigation using received signal strength. *IEEE Transactions on Wireless Communications*, 14(3):1689–1702, 2015.

[31] R. Ye, S. Redfield, and H. Liu. High-precision indoor UWB localization: Technical challenges and method. In *2010 IEEE International Conference on Ultra-Wideband*, volume 2, pages 1–4, 2010.