

Final Report

ROBOTIC ULTRASOUND ASSISTANCE VIA
HAND-OVER-HAND CONTROL

EN 601.656 Computer Integrated Surgery II

Kevin Gilboy
Computer Engineering M.S.E. Student
kevingilboy@jhu.edu

Table of Contents

Table of Contents	i
1 Introduction	1
1.A Clinical Motivation	1
1.B Prior Work	1
1.C Goal	2
2 Experimental Setup	3
2.A Hardware	3
2.B Software	3
3 Technical Approach	4
3.A Robot Interfacing	5
3.B Force Sensor Interfacing	5
3.B.1 Hand Forces: Robotiq FT-150	5
3.B.2 Probe Contact Forces: Variense FSE103 and Honeywell Model 31	6
3.C Force Aligning	10
3.D Gravity Compensation	11
3.E Kalman Filtering	12
3.F Admittance Control	14
3.F.1 In-Air Admittance Control	14
3.F.2 Compensating For Probe Contact Forces	15
3.F.3 Demoing While Using a Broken Probe Force Sensor	16
4 Results	16
4.A Probe Contact Force Using Neural Networks	16
4.B Gravity Compensation	19
4.C Kalman Filtering	21
4.D In-Air Admittance Control	23
5 Discussion	24
5.A Probe Contact Force Using Neural Networks	24
5.B Gravity Compensation	25
5.C Kalman Filtering	25
5.D In-Air Admittance Control	26
6 Progress Evaluation	26
6.A Dependencies	26
6.B Activities and Deliverables	27
6.C Schedule Adherence	28
6.D Reflection	29
7 Conclusion	30
7.A Significance	30

7.B Moving Forward	30
7.C Final Word	30
References	31

1. Introduction

1.A Clinical Motivation

This project is motivated by the fact that ultrasound (US) guided procedures typically require a sonographer to hold an US probe against a patient in static, contorted positions for long periods of time while also applying large forces [1]. As a result, 63%-91% of sonographers develop occupation-related musculoskeletal disorders compared to only about 13%-22% of the general population [2].

The vision of this work, as well as the previous works discussed in the next section, is to provide sonographers with “power-steering” via a hand-guidable robot that they can maneuver to a point-of-interest and then release, having the robot do all the strenuous holding on their behalf.

1.B Prior Work

Prior work in robotic US assistance has been performed by numerous individuals at JHU, with the most notable contributions coming from Rodolfo Finocchi and Ting-Yun (Angel) Fang whose work serves as a starting point for this project. Finocchi was an original developer of the first robotic US assistance prototype, in which he used MATLAB code to add admittance control to a UR5 robot using dual-force sensors (to decouple forces applied to a patient) in a custom housing. His main contributions included the custom US probe and dual-force sensor housing, algorithms for admittance control and contact force control, use of a 1€ filter to smooth force/torque (F/T) inputs, use of a sigmoidal force-to-velocity conversion function, and user-study verification through grip force measured via FlexiForce film [3, 4]. Finocchi’s setup is shown in Fig. 1 below.




Figure 1: The experimental setup developed and published by Finocchi [3, 4]. The custom probe and dual-force sensing housing can be seen attached to the UR5 robot.

As an application of Finocchi’s work, Zhang used the developed cooperatively controlled robotic system in combination with virtual fixtures to implement synthetic tracked aperture US (STRATUS) imaging [5]. Additionally, Finocchi’s robotics work was extended by Fang

who iterated upon the custom housing design to make it more compact and validated the result through applied force reduction, stability of contact force, and stability of US images [6]. Fang’s improved setup is shown in Fig. 2 below.




Figure 2: The experimental setup developed and published by Fang [6]. The improved end-effector housing can be seen attached to the UR5 robot.

While all of these works were successful in relieving grip strain while maintaining force profiles and capturing quality US images, multiple users (namely Dr. Russell Taylor and Dr. Emad Boctor) had reported that their prototypes lacked the “transparency” of power-steering necessary for clinical usage.

1.C Goal

In the previous section, it was shown that previous attempts at cooperatively-controlled robotic US with dual-force sensing were successful in many applications such as maintaining force profiles, enforcing virtual fixtures, and relieving grip strain, but lacked the overall motion transparency to be practical in a clinical setting. The goal of this work is to improve upon the previous robotic US assist prototypes and create a more transparent power-steering, cooperative-control experience for sonographers. While the same sensors, housing, and robot will be used, it is believed that additional algorithms can be developed to help smooth the commanded robot motion.

If successful and validated, this work will be an important progression toward mitigating sonographers’ susceptibility to work-related musculoskeletal disorders. It also has consequences for all procedures under the umbrella of robotic US, as the control algorithms developed for this work will underlay and improve all applications built on top of it. Some examples include enforcing virtual fixtures for synthetic aperture procedure, imaging with respiratory gating, replicating a position/force profile for repeatable biopsy procedures, and conducting co-robotic US tomography scans.

2. Experimental Setup

2.A Hardware

The hardware setup was as shown in Fig. 3 below. A desktop computer was used to communicate with a 6-DoF UR5 robot using TCP/IP on a private LAN. Through this connection, the computer received position, velocity, and force information while transmitting commanded velocities. The dual-force sensing, US probe wielding end-effector developed by Fang [6] was attached to the robot. The system used “dual-force” sensing since in addition to the 6-DoF Robotiq FT-150 F/T sensor that measured applied hand forces, a second force sensor was used to discern forces applied by the probe against the patient or phantom. Two force sensors were available for this task: a 1-DoF Honeywell Model 31 mid-range load cell and a 3-DoF Variense FSE103 force sensor. While issues were found with both sensors, the latter was primarily used in this work since the former produced very unusable results (more described in Section 3.B.2).




Figure 3: The experimental setup for this work, which includes a computer commanding a UR5 robot over TCP/IP as well as probe-holding end-effector developed by Fang [6] (figure inside the black border is from Fang, 2017). While a 1-DoF load cell is shown in the diagram, a 3-DoF force sensor was also used in its place since the 1-DoF sensor produced unusable results.

2.B Software

As mentioned in Section 1.B, Finocchi [3, 4] and Fang [6] implemented their code in MATLAB and used client software running on the UR5 to relay F/T values to the computer. While they still achieved an adequate result, using an interpreted language such as MATLAB and running unnecessary client-side code introduces latency and overhead which is detrimental to the user experience of any real-time system.

In this work, C++ was used in combination with the open-source [CISST/SAW](#) libraries to get data from, and command the UR5 without any client-side code. A simplified diagram is shown in Fig. 4.

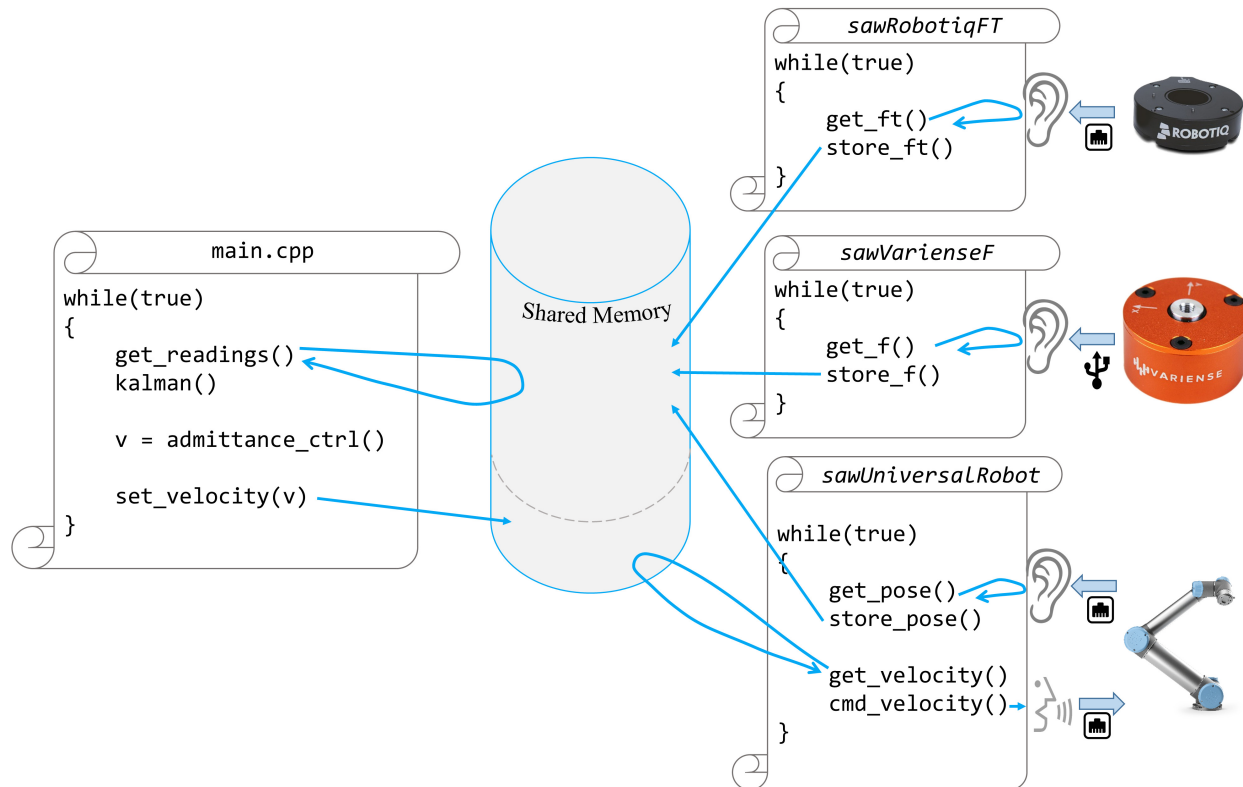


Figure 4: A simplified code flow diagram showing the asynchronous component listeners and `main.cpp` which performed the admittance control, filtering, and robot commanding.

As shown, there were three SAW components listening for data from the robot and F/T sensors respectively and storing them in objects accessible by `main.cpp` through CISST state table APIs. While the `sawUniversalRobot` component was previously written by the SMARTS lab, the `sawRobotiqFT` and `sawVarienseF` components were created specifically for this project by the author. The main script, in addition to performing component initialization, was essentially an infinite loop of fetching readings, filtering, and commanding a velocity to the UR5 SAW component. It is worth noting that the CISST/SAW libraries have native support for accessing shared data in a way that prevents race conditions, which is very useful since this program relies on asynchronous, multitask execution and is therefore prone to data corruption.

3. Technical Approach

There are many ways to present the technical aspects of this work. In this section, the technical approach for the complete hand-over-hand control algorithm (from hardware interfacing to admittance control) will be described sequentially in subsections. A graphical overview

of the dataflow is shown in Fig. 5, with each triangle representing a signal-manipulating functional component that will be described in a subsequent subsection.

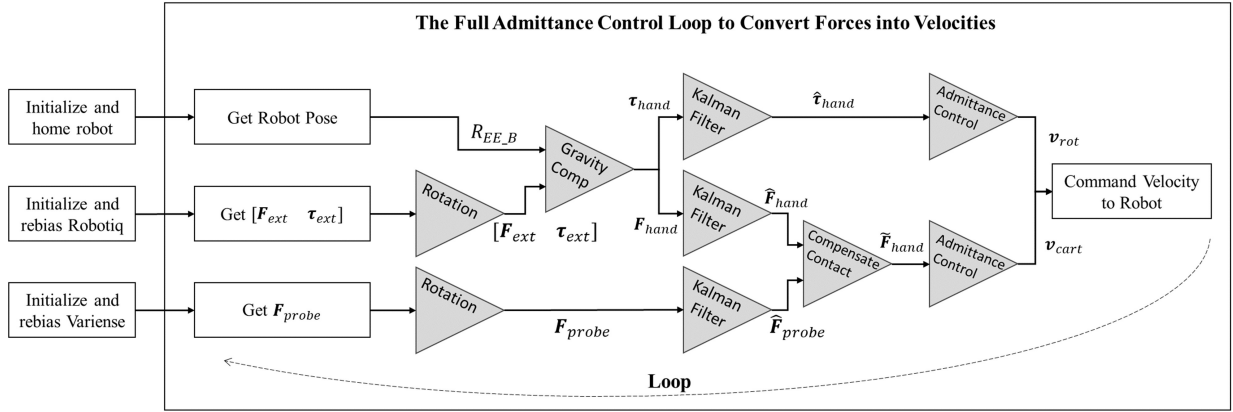


Figure 5: The overall dataflow and passage of signals for the implemented admittance control loop. After initialization, the loop includes frame transformations, gravity compensation, Kalman filtering, contact compensation, and admittance control. Each of these aspects are described in detail later in this section.

While plots and diagrams will be used to describe design decisions, results of these individual technical aspects are reserved for Section 4.

3.A Robot Interfacing

Interfacing with the UR5 was performed via the CISST `sawUniversalRobot` component (pulled on Jan. 2019) which communicated with the robot via TCP over its static IP address 192.168.1.142 and preconfigured communication port 30004. The CISST component was used to get the robot’s status and command linear/rotational Cartesian velocities deduced by the eventual admittance control algorithms. Using this component greatly simplified the robot control implementation since knowledge of the forward and inverse kinematics could be abstracted away behind a neat and reliable API.

3.B Force Sensor Interfacing

3.B.1 Hand Forces: Robotiq FT-150

The Robotiq 6-DoF F/T sensor was used for measuring hand forces \mathbf{F}_{hand} and torques $\boldsymbol{\tau}_{hand}$, and had a digital output that was natively compatible with the UR5. The Robotiq sensor was wired to the UR5 main controller where the values were implicitly assembled into TCP packets and forwarded to port 63351 on the UR5’s static IP address. A custom `sawRobotiqForceSensor` module was created to 1) rebias the sensor, 2) asynchronously listen for data packets on the Robotiq data port, and 3) parse these values into a state table upon reception which could then be asynchronously and safely accessed by the main program through typical CISST APIs. Rebiasing took place over two steps at the start of the program once the robot was navigated to a home position that had the end-effector pointing vertically

down. First, a special rebias command was sent over TCP/IP to the Robotiq which zeroed out the sensor at its instantaneous value. This could be a problem if the sensor was rebiased at the peak of a noise spike, so second a rebias vector was calculated by averaging 500 F/T samples together to create a bias correction factor, which could then be subtracted from any future force reading to yield the corrected reading. This two step procedure ensured robustness to noise.

3.B.2 Probe Contact Forces: Variense FSE103 and Honeywell Model 31

There were two force sensors that were configured for measuring probe contact forces \mathbf{F}_{probe} in this project, each with 3D-printed housing fabricated by previous researchers. The second force sensor was needed to discern hand forces from “equal and opposite” forces applied by the phantom/patient onto the probe when contact was made. Without discerning and compensating, the probe would bounce off the contacted surface due to the opposite forces. This relevance is further described in Section 3.F.2. While the eventual system used a 3-DoF Variense force sensor, it is useful to describe some of the hurdles experienced along the way, as well as the work performed with the alternative 1-DoF Honeywell Model 31 mid-range load cell.

Variense 3-DoF Force Sensor The Variense 3-DoF force sensor had a serial USB output with a data rate of 50Hz that could be read from `/dev/ttyACM0` on a Linux machine. Although the online documentation was incorrect, the serial packet structure was determined through functional testing to be the following in MSB \rightarrow LSB order:

```

1 struct VarienseDataPacket
2 {   byte    0x00 ,
3     uint32  rel_timestamp_ms , // ~20ms between readings
4     float32 Fx ,
5     float32 Fy ,
6     float32 Fz ,
7     byte    0xFF
8 };

```

A custom `sawVarienseForceSensor` module was created to 1) rebias the sensor, 2) asynchronously listen for new data packets over the Serial connection, and 3) parse these values into a state table which could then be asynchronously and safely accessed by the main program through typical CIST APIs. Rebiasing took place at the start of the program once the robot was navigated to a home position that had the end-effector pointing vertically down. Similar to the Robotiq rebiasing, a rebias vector was calculated by averaging 500 samples together, which could then be subtracted from any future force reading to yield the corrected reading.

While the interfacing worked without trouble, nonlinear force redirection introduced by compliance and hysteresis in the force sensor housing proved to be difficult to overcome. The observed behavior is plotted in Fig. 6 when forces were applied directly to the probe. In an ideal scenario, the observed \mathbf{F}_{hand} should equal \mathbf{F}_{probe} since forces applied to the probe and measured by the Variense force sensor are also felt by the upstream Robotiq force sensor. However, it was observed that the forces were quite different in their component strengths and magnitudes.

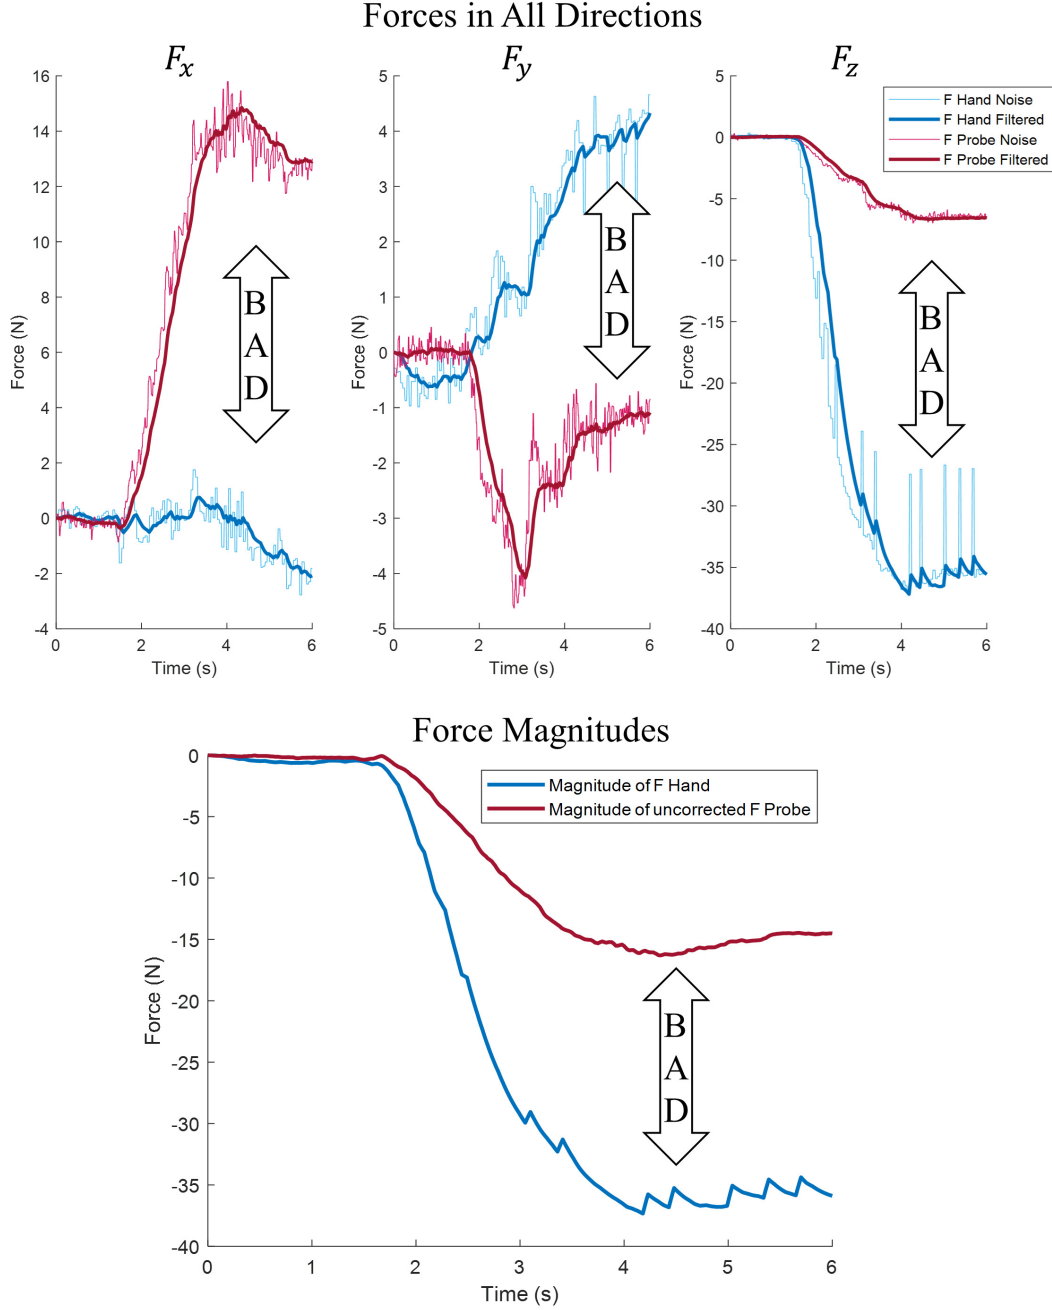


Figure 6: Multiple plots showing the incongruity between \mathbf{F}_{hand} measured by the Robotiq and \mathbf{F}_{probe} measured by the Variense when forces were applied directly to the probe. In this scenario, ideally $\mathbf{F}_{hand} == \mathbf{F}_{probe}$ since forces applied to the Variense should be equivalently felt by the Robotiq. Since the forces are redirected in a stochastic way, it was very hard to compensate \mathbf{F}_{hand} for probe forces applied to the patient.

An attempt at a software solution was made by training a neural network to characterize and compensate for these nonlinearities. The model, shown in Fig. 7, consisted of three multiple-input, single-output networks to deduce the corrected $\hat{\mathbf{F}}_{probe}$ from the original \mathbf{F}_{probe} . Each network consisted of one hidden layer of 20 neurons with a sigmoidal activation function,

and one output layer.

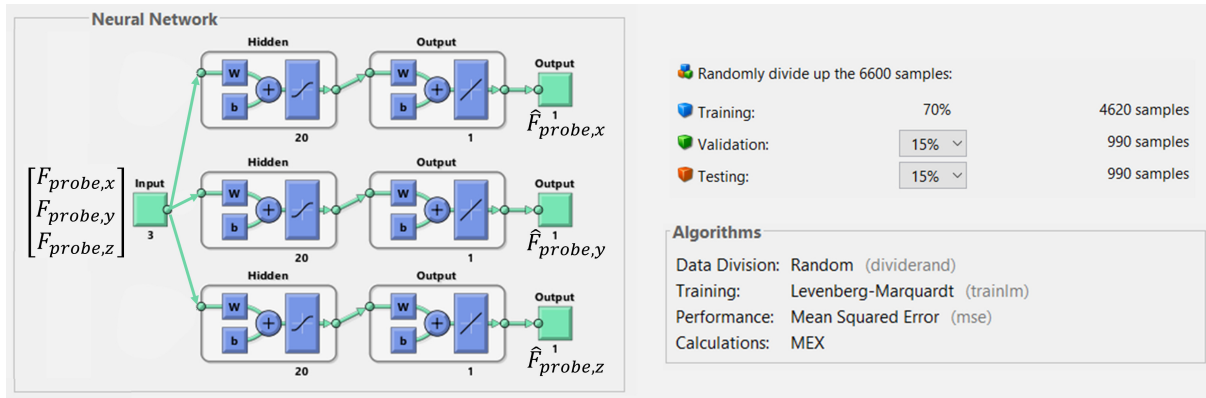


Figure 7: The architecture of the neural network used in an attempt to predict true probe forces given the probe force measurement.

While the neural net did make improvements to the probe measurement, most significantly in equalizing the magnitude of the force sensors, it was deemed to be too ineffective for practical usage. Most importantly, its seemingly random spikes during usage and changes of characterization between usage were very unrealistic for a clinical setting. This is discussed much further in Section 5.A.

Honeywell 1-DoF Load Cell The Honeywell load cell had an analog output, for which an INA125p microelectronic amplifier and analog-to-digital conversion circuit was created. The schematic is shown in Fig. 8.

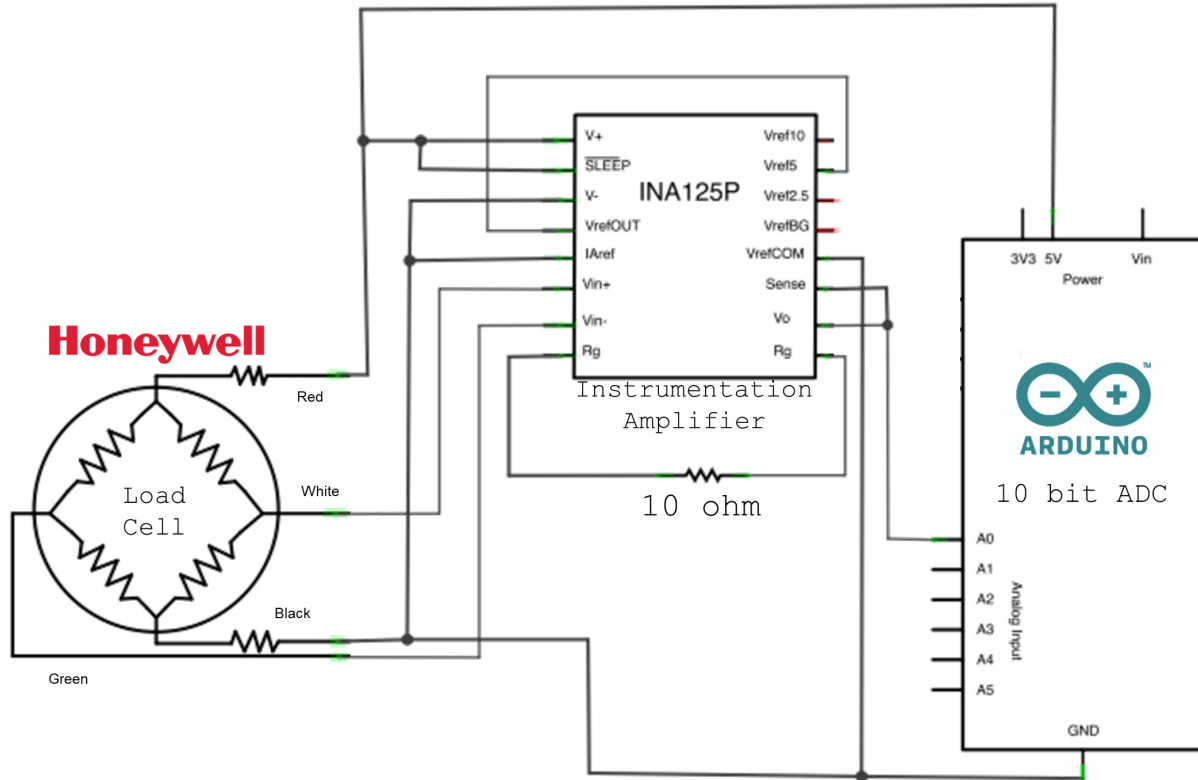


Figure 8: The amplifier and analog-to-digital conversion circuit used to convert load cell excitation voltages into a serial string that could then be read by a custom CISST component. An Arduino with a built-in 10-bit ADC and amplifier wired into a protoboard shield was used since it could be prototyped and evaluated without much time cost.

An Arduino Uno with an onboard 10-bit analog-to-digital converter was used to convert the analog readings into a serial string that could then be read by a serial CISST component such as the custom `sawVarienseForceSensor` created earlier. Unfortunately, the sensor was determined through testing to be in a constant state of saturation, perhaps from previous improper storage, and therefore unusable for this experiment. This could be deduced by observing the load cell's analog voltage while applying different forces to the load cell, the results of which are shown visually in Fig. 9. When compression forces were applied, the load cell experienced a noisy deflection of $+0.005V$ which is a one bit change at a 10-bit resolution for a maximum analog signal of $5V$ (since $\frac{5V}{2^{10} \text{ levels}} \approx 0.0049 \frac{Volts}{Bit}$). When tensional forces were applied, the load cell experienced a deflection of $-0.5V$. Since the load cell should not register tensional forces and only experienced an insignificant deflection for compression forces, it could logically be determined that the load cell was compressed together even when unperturbed, leaving it in a state of saturation. Therefore, this sensor was considered to be broken beyond repair and not used further in this work.

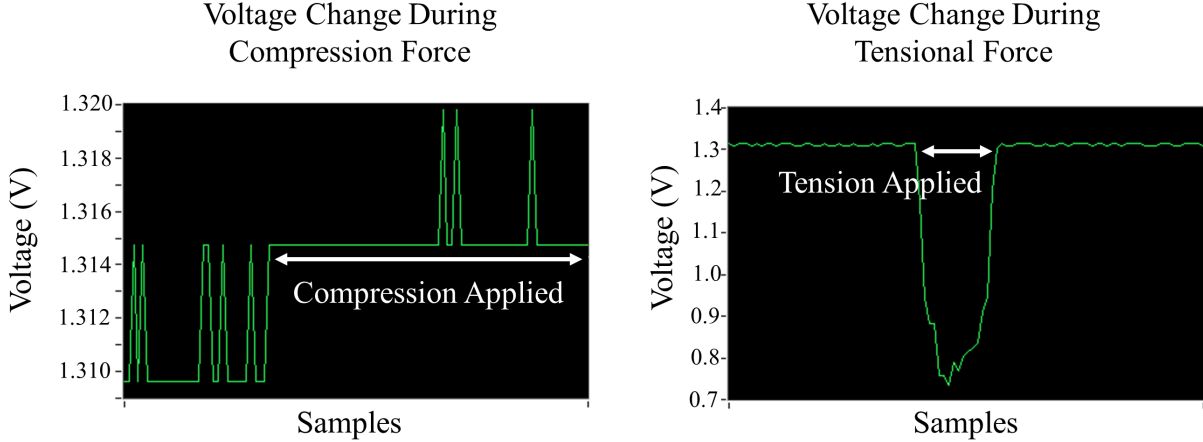


Figure 9: The measured voltage change when compression and tensional forces were applied to the Honeywell load cell. Since compression forces only caused a 1-bit deflection and tensional forces were greatly registered by the load cell, it could be concluded that the cell was in a state of constant saturation due to improper storage and therefore unfixable/unusable.

3.C Force Aligning

Due to the orientation of the force sensors with respect to the robot, rotation matrices were used to reorient the observed forces into the robot base frame. While this could be solved by a least-squares method, all force sensors were rigidly fixed to the robot at convenient angles so the matrices could be determined by induction.

- Rotation of Robotiq 6-DoF force torque sensor: This sensor was oriented with the z-axis coaxial to the robot base frame z-axis and the x/y axes rotated $-\frac{\pi}{2}$ about the z-axis. For the base frame to the Robotiq, this results in the ZYX Euler angles $(-\frac{\pi}{2}, 0, 0)$ and rotation matrix

$$R = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

- Rotation of Variense 3-DoF force sensor: This sensor was oriented “sideways” with the z-axis parallel to the xy-plane of the robot base frame and the x-axis coaxial with the robot base frame z-axis. For the base frame to the Variense, this results in the ZYX Euler angles $(\frac{\pi}{2}, -\frac{\pi}{2}, 0)$ and rotation matrix

$$R = \begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & -1 \\ 1 & 0 & 0 \end{bmatrix} \quad (2)$$

- Rotation of Honeywell 1-DoF load cell: This sensor was oriented with the z-axis pointing opposite of the robot base frame z-axis. Since this sensor only had 1-DoF, multiple solutions exist. For the base frame to the Honeywell, one possible solution would be

the ZYX Euler angles $(0, \pi, 0)$ and rotation matrix

$$R = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{bmatrix} \quad (3)$$

3.D Gravity Compensation

Gravity (tool weight) compensation was performed on each reading from the Robotiq 6-DoF F/T sensor based on the robot's pose at that instant. This is because the reading of external forces from the Robotiq, $[\mathbf{F}_{ext} \ \boldsymbol{\tau}_{ext}]$, is a summation of the forces caused by the operator's hand $[\mathbf{F}_{hand} \ \boldsymbol{\tau}_{hand}]$ and the forces induced from gravity, $[\mathbf{F}_{grav} \ \boldsymbol{\tau}_{grav}]$. Our goal is to solve for $[\mathbf{F}_{hand} \ \boldsymbol{\tau}_{hand}]$ through the measurement of $[\mathbf{F}_{ext} \ \boldsymbol{\tau}_{ext}]$ and analytical deduction of $[\mathbf{F}_{grav} \ \boldsymbol{\tau}_{grav}]$ based on robot pose:

$$\begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{hand} = \begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{ext} - \begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{grav} \quad (4)$$

Several frames were defined as shown in Fig. 10, which allow $[\mathbf{F}_{grav} \ \boldsymbol{\tau}_{grav}]$ in equation (4) to be written in terms of solvable constants.

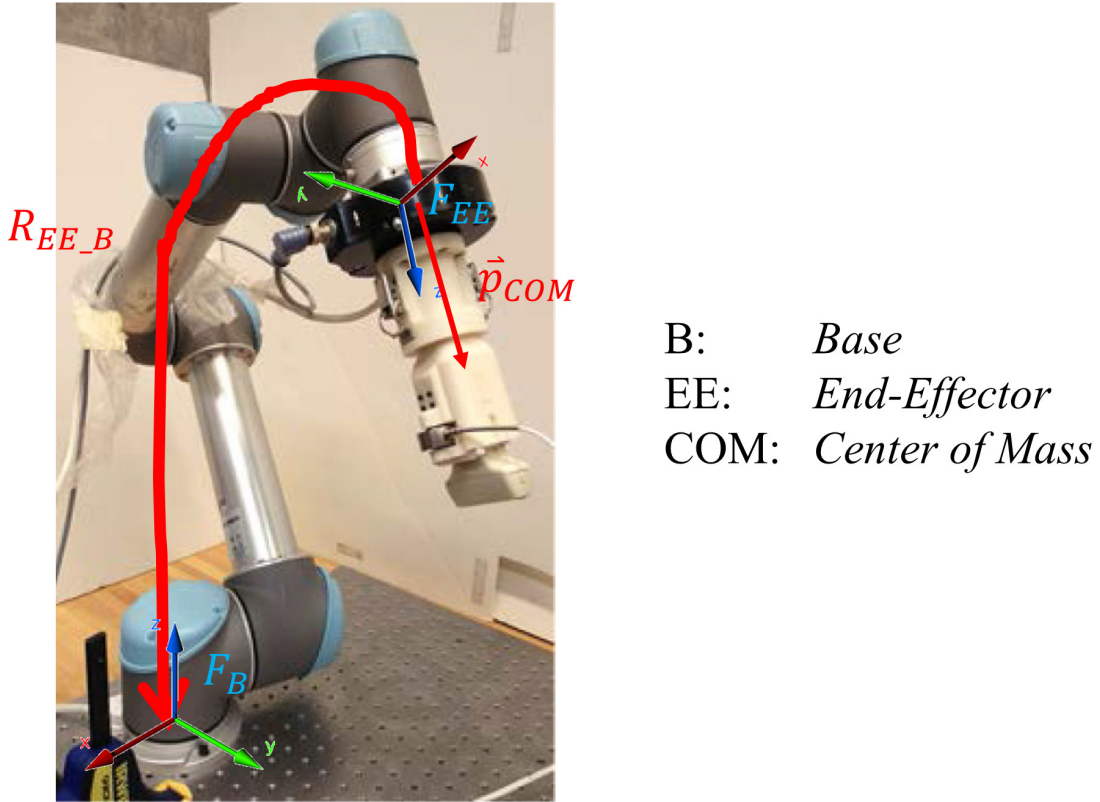


Figure 10: The relevant frames (blue) and rotations/translations (red) used for gravity compensation.

Given these frames and the knowledge of the procedure used by [7], equation (4) can be rewritten as

$$\begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{hand} = \begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{ext} - \left(\begin{bmatrix} R_{EE.B} & 0 \\ -R_{EE.B} \cdot sk(\vec{p}_{COM}) & R_{EE.B} \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ -mg \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -mg \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \quad (5)$$

where the second addition of the $-mg$ is meant to “undo” the sensor rebiasing and add in the expected amount of gravity contribution. Compensation requires knowledge or deduction of the weight contribution of the US probe and housing $-mg$ when in the vertical position, as well as the x and y center-of-mass coordinates for the probe p_{COM} . Both of these were solved analytically as described in [7] using 1000 samples of data for 32 known static poses around the workspace, all collected by an automated batch test script. Once the samples for each pose were averaged together, since there were no hand forces applied during data collection, equation (5) could be simplified to

$$\left(\begin{bmatrix} R_{EE.B} & 0 \\ -R_{EE.B} \cdot sk(\vec{p}_{COM}) & R_{EE.B} \end{bmatrix} + I_{6 \times 6} \right) \begin{bmatrix} 0 \\ 0 \\ -mg \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{ext} \quad (6)$$

which is in $Ax = b$ form and can be solved in a least-squares sense for $-mg$. Once $-mg$ is known, the x and y coordinates of \vec{p}_{COM} can subsequently be solved since equation (6) can be rearranged to

$$R_{EE.B} \cdot sk(\vec{p}_{COM}) \begin{bmatrix} 0 \\ 0 \\ -mg \end{bmatrix} = \begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{ext} \quad (7)$$

Using the definition of a skew-symmetric matrix, this can be again rearranged to

$$\left(R_{EE.B} \cdot \begin{bmatrix} -mg & 0 & 0 \\ 0 & -mg & 0 \\ 0 & 0 & -mg \end{bmatrix} \right) \begin{bmatrix} p_{COM,y} \\ -p_{COM,x} \\ 0 \end{bmatrix} = \begin{bmatrix} \mathbf{F} \\ \boldsymbol{\tau} \end{bmatrix}_{ext} \quad (8)$$

which is in $Ax = b$ form that can be solved in a least-squares sense for $p_{COM,x}$ and $p_{COM,y}$.

Once $-mg$, $p_{COM,x}$, and $p_{COM,y}$ were analytically determined, they could be plugged into (5) as part of the admittance control loop to convert a measured $\begin{bmatrix} \mathbf{F}_{ext} & \boldsymbol{\tau}_{ext} \end{bmatrix}$ into the desired $\begin{bmatrix} \mathbf{F}_{hand} & \boldsymbol{\tau}_{hand} \end{bmatrix}$.

3.E Kalman Filtering

The force sensors used in this project experienced a noise range of up to $2N$ ($\pm 1N$) and were the main “rate-limiting factor” of the admittance control loop since they sent packets at a

speed up to 6.25X slower than the robot could be commanded. Therefore, Kalman filtering was employed on readings from both force sensors to 1) smooth noise, 2) predict future movements through observation, and 3) infer inter-packet force readings as shown below.

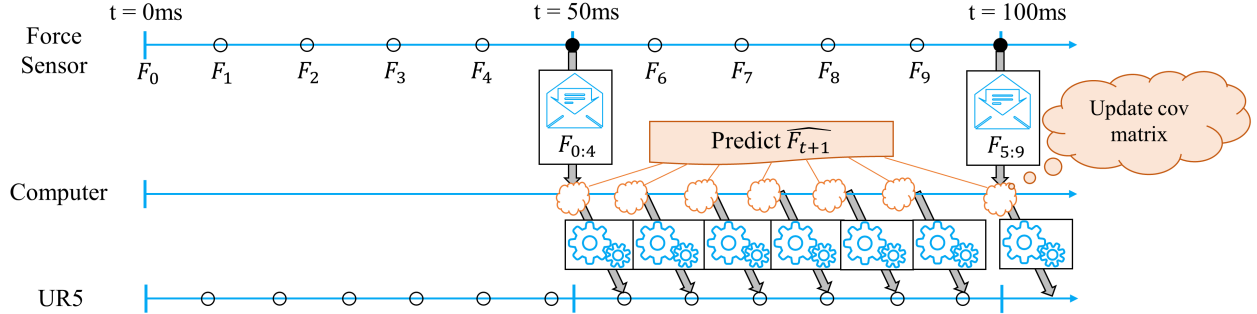


Figure 11: A visual display of how Kalman filtering can be used to “upsample” low-rate ground-truth force data so that the robot can be commanded as fast as possible. Between force data packets, the Kalman filter infers and predicts future force values allowing for fast and smooth commanding of the UR5 robot to enhance motion transparency.

Three separate Kalman filters were instantiated and used due to noise differences and in the sensors: one for filtering hand forces, one for filtering hand torques, and one for filtering probe forces. Namely, this produced $\hat{\mathbf{F}}_{hand}$, $\hat{\boldsymbol{\tau}}_{hand}$, and $\hat{\mathbf{F}}_{probe}$. In doing so, the motion transparency experienced by the user was expected to greatly improve.

Kalman filtering was performed using a sequential procedure that both predicts subsequent values and updates internal covariance matrices to improve future predictions. The algorithm was built on the assumption of approximately constant acceleration between samples, and takes an input vector $m \in \mathbb{R}^{9 \times 1}$ of nine values: three for the current force, three for the first derivative of the force (sometimes called “velocity”), and three for the second derivative of the force (sometimes called “acceleration”). Other notable variables include the manually tuned $R \in \mathbb{R}^{9 \times 9}$ sensor confidence covariance and $Q \in \mathbb{R}^{9 \times 9}$ action uncertainty covariance, measurement selector $H = I_{6 \times 6}$, persistent state estimate $x \in \mathbb{R}^{6 \times 1}$, and transition matrix

$$\begin{bmatrix} 1 & 0 & 0 & dt & 0 & 0 & \frac{1}{2}dt^2 & 0 & 0 \\ 0 & 1 & 0 & 0 & dt & 0 & 0 & \frac{1}{2}dt^2 & 0 \\ 0 & 0 & 1 & 0 & 0 & dt & 0 & 0 & \frac{1}{2}dt^2 \\ 0 & 0 & 0 & 1 & 0 & 0 & dt & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & dt & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & dt \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (9)$$

with $dt = 0.01$ for the 100Hz that the filter will be run at. The Kalman algorithm can be summarized as the below steps.

Prediction

$$\begin{aligned} x[t] &= Ax[t - 1] && \text{(State Estimate)} \\ P &= (APA^T) + Q && \text{(Predicted Error Covariance)} \end{aligned}$$

Update

$$\begin{aligned} S &= HPH^T + R && \text{(Pre-fit Residual Covariance)} \\ K &= PH^T S^{-1} && \text{(Optimal Gain)} \\ y &= m - Hx[t] && \text{(Pre-fit Residual)} \\ x[t] &= x[t] + Ky && \text{(Updated State Estimate)} \\ P &= (I - KH)P && \text{(Update Estimate Covariance)} \end{aligned}$$

After the above algorithm was implemented, the Q and R covariances were manually tuned first in MATLAB simulation, then tweaked based on qualitative user feedback and graphical evaluation of rise time, overshoot, and phase-lag (felt as “inertia”).

3.F Admittance Control

3.F.1 In-Air Admittance Control

Admittance control, ignoring probe forces, was implemented using linear force-to-velocity gains to convert hand forces and torques measured by the Robotiq force sensor into Cartesian velocities commanded to the robot. The general equation using α for slope and η for expected noise is:

$$\dot{x} = \begin{cases} 0 & |F_{hand}| \leq \eta \\ \text{sgn}(F_{hand}) \cdot \alpha (|F_{hand}| - \eta) & |F_{hand}| > \eta \end{cases} \quad (10)$$

This function has a deadband zone when the measured force is within the noise threshold $[-\eta, \eta]$ to attenuate force noise leading to vibration, but otherwise increases linearly with force. A sample plot is shown below.

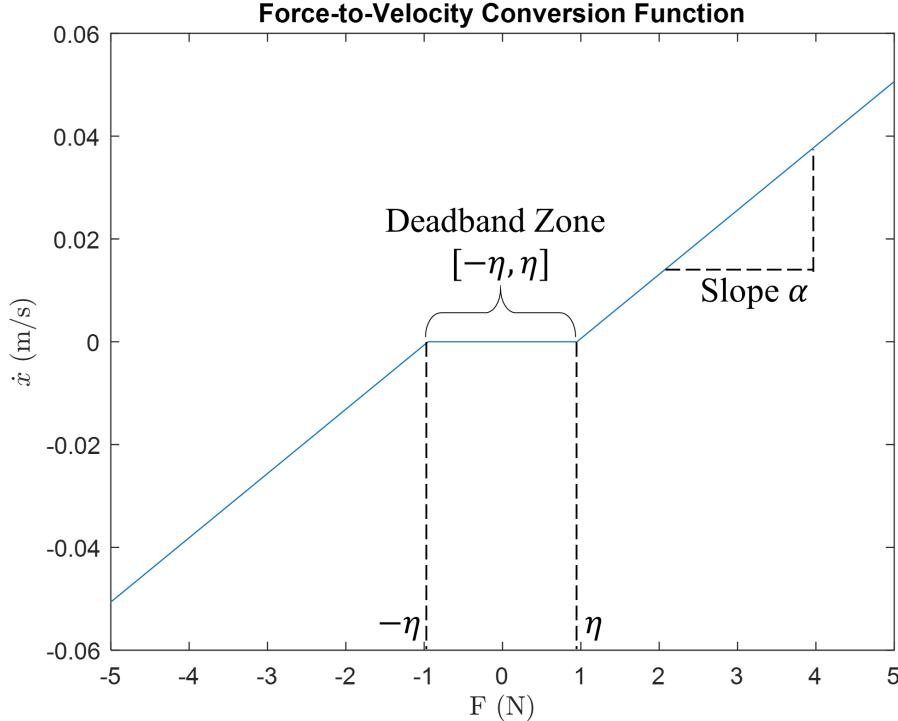


Figure 12: A graphical view of the piecewise linear admittance control function to convert forces into commandable velocities. The deadband zone helped attenuate the translation of force noise into vibration at the cost of less responsive behavior, and the slope could be tuned to speed up or down the robot response.

The control constants were tweaked through qualitative feedback to be $\alpha = 0.0125$, $\eta = 0.95$ for the x/y-directions and $\alpha = 0.0125$, $\eta = 0.25$ for the z-direction. The z-direction was given a smaller noise threshold since the Robotiq force sensor has less noise in the z-direction (shown in Section 4.C), and it makes the system more responsive for z-direction palpation during scanning.

The use of linear admittance control gains was believed to be an upgrade from the sigmoidal gains used by [3, 4]. After evaluating plots of motion for both admittance control schemes and discussing, it was determined that using sigmoidal gains introduced nonlinearities into the Kalman filter which was not designed to model nonlinear behavior. While the proposed admittance control function is only piecewise linear and non-differentiable, it still resulted in better control than the sigmoidal function as shown in Section 4.D. It should be noted that since [3, 4] used frequency-domain filtering instead of Kalman filtering, their work may not have been as affected by the nonlinearities of the sigmoidal curve.

3.F.2 Compensating For Probe Contact Forces

A second force sensor, as stated before, was necessary for discerning probe forces from hand forces when a scan is occurring. Without discerning, the “equal-and-opposite” forces exerted by the elastic phantom/patient onto the probe will be detected as part of \mathbf{F}_{hand} by the Robotiq sensor and cause the robot to “bounce” off the scanning surface since the force is in the opposite direction. Therefore, what has been referred to as \mathbf{F}_{hand} all along is

actually $\mathbf{F}_{hand} = \tilde{\mathbf{F}}_{hand} + \mathbf{F}_{probe}$, where $\tilde{\mathbf{F}}_{hand}$ is the true force exerted by the user's hand. In the in-air case, $\mathbf{F}_{probe} = 0$ meaning $\mathbf{F}_{hand} = \tilde{\mathbf{F}}_{hand}$ which is why we could use equation (10) for admittance control previously. When scanning, however, \mathbf{F}_{probe} is nonzero due to contact with the patient and therefore the admittance control becomes

$$\dot{x} = \begin{cases} 0 & |\tilde{F}_{hand}| \leq \eta \\ \text{sgn}(\tilde{F}_{hand}) \cdot \alpha (|\tilde{F}_{hand}| - \eta) & |\tilde{F}_{hand}| > \eta \end{cases} \quad (11)$$

or equivalently

$$\dot{x} = \begin{cases} 0 & |F_{hand} - F_{probe}| \leq \eta \\ \text{sgn}(F_{hand} - F_{probe}) \cdot \alpha (|F_{hand} - F_{probe}| - \eta) & |F_{hand} - F_{probe}| > \eta \end{cases} \quad (12)$$

3.F.3 Demoing While Using a Broken Probe Force Sensor

As mentioned in Section 3.B.2, neither attained probe force sensor worked sufficiently enough to implement the ideal probe force compensation algorithm proposed in the previous section. While this is an issue that needs addressed for future work, an approach was taken to still allow the project functionality to be demonstrated at the class poster session. The approach used the force magnitude from the 3-DoF Variense force sensor, and upon feeling a measured contact force greater than some threshold, would apply a virtual fixture directly above the probe to prevent it from bouncing upward while allowing further uninhibited movement downward. The virtual fixture, in other words, would effectively enforce contact between the the probe and surface as long as some contact force was registered. When the probe was dragged off of the scanned surface thereby relieving the contact force, the virtual fixture would no longer be enforced allowing the probe to move upward in the z-direction once again. While the implementation was limited and used a horizontal virtual fixture while assuming a horizontal scanning surface, it was temporarily acceptable for showing functionality.


4. Results

In addition to the overall result of this project, the results of the individual technical approaches described in Section 3 are equally important to outline.


4.A Probe Contact Force Using Neural Networks

As mentioned in Section 3.B.2, the 1-DoF Honeywell load cell was not used in this experiment since it was determined to be broken. The main result presented here was for using a neural network to characterize the nonlinearities of the 3-DoF Variense force sensor. The results of the neural net training are shown below. The correlation coefficient on the validation data in the x-, y-, and z-directions respectively were 0.945, 0.981, and 0.998.

Correction of Probe Forces in X-Direction



Correction of Probe Forces in Y-Direction



Correction of Probe Forces in Z-Direction




Figure 13: The results of training the neural network to characterize true probe forces based on those measured by the Variense. As seen, all networks converged and had strong correlation coefficients for the validation and test datasets.

The performance of the network on two test datasets are shown below, the former showing successful force characterization by the network and the latter showing an unsuccessful force characterization. In these cases, force was applied to the probe in the absence of any applied F_{hand} meaning that in an ideal scenario where the sensor worked properly, the red line representing F_{probe} would overlap with the blue line representing F_{hand} so that the true \hat{F}_{hand} could be calculated as $F_{hand} - F_{probe} = 0$. The gold line shows the neural network prediction of the probe forces \hat{F}_{probe} calculated from the measured F_{probe} , which in an ideal scenario where the network correctly characterized the nonlinear force sensor readings would also overlap with the blue line. These datasets were taken two days apart, the significance of which will be discussed further in Section 5.A.

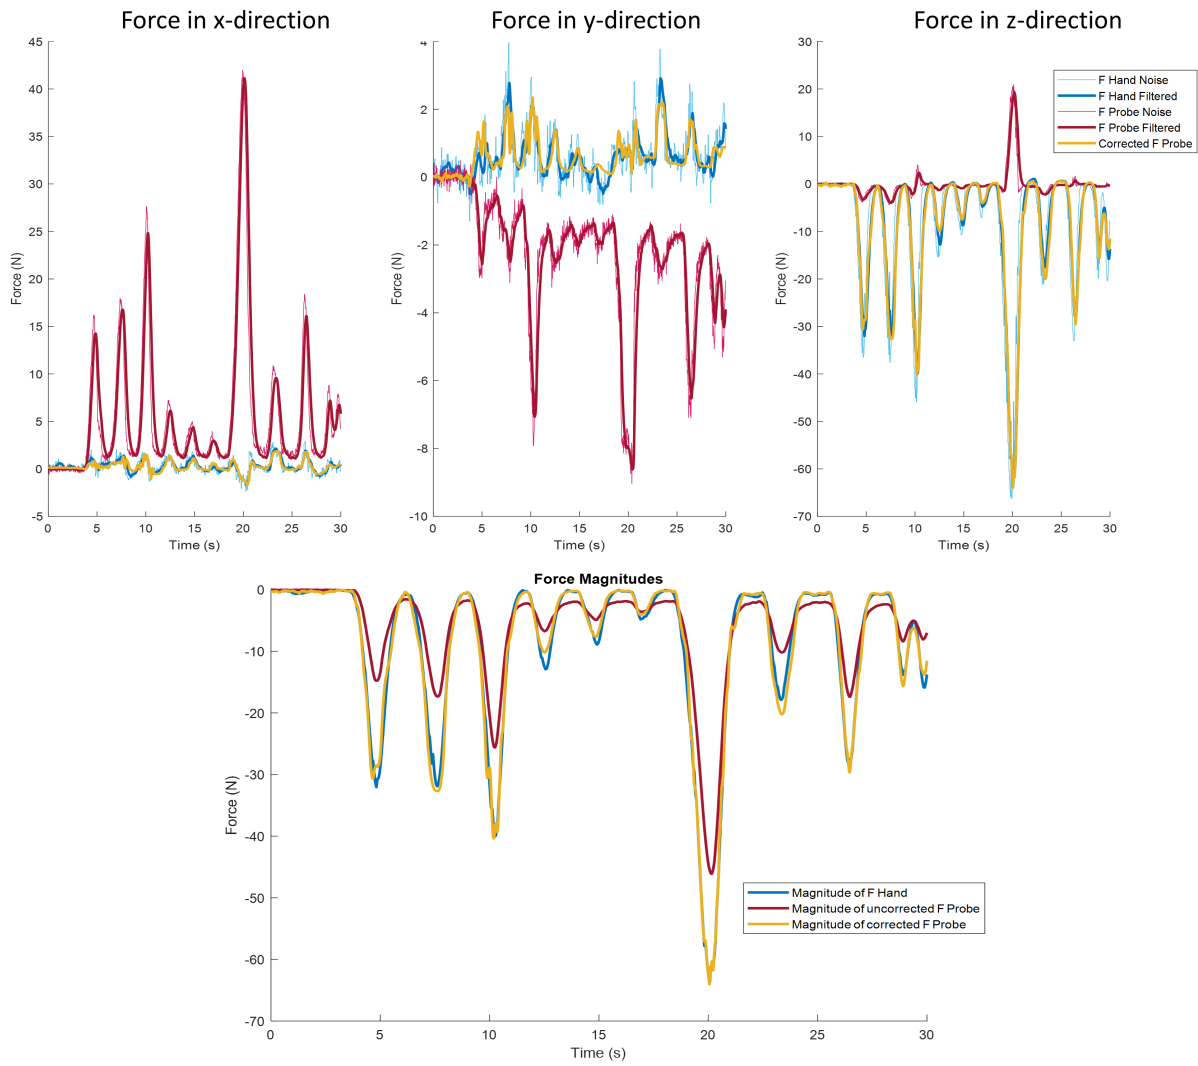


Figure 14: Dataset one of \mathbf{F}_{hand} , \mathbf{F}_{probe} , and the corrected $\hat{\mathbf{F}}_{probe}$ from the neural network, collected on the same day as the network was trained.

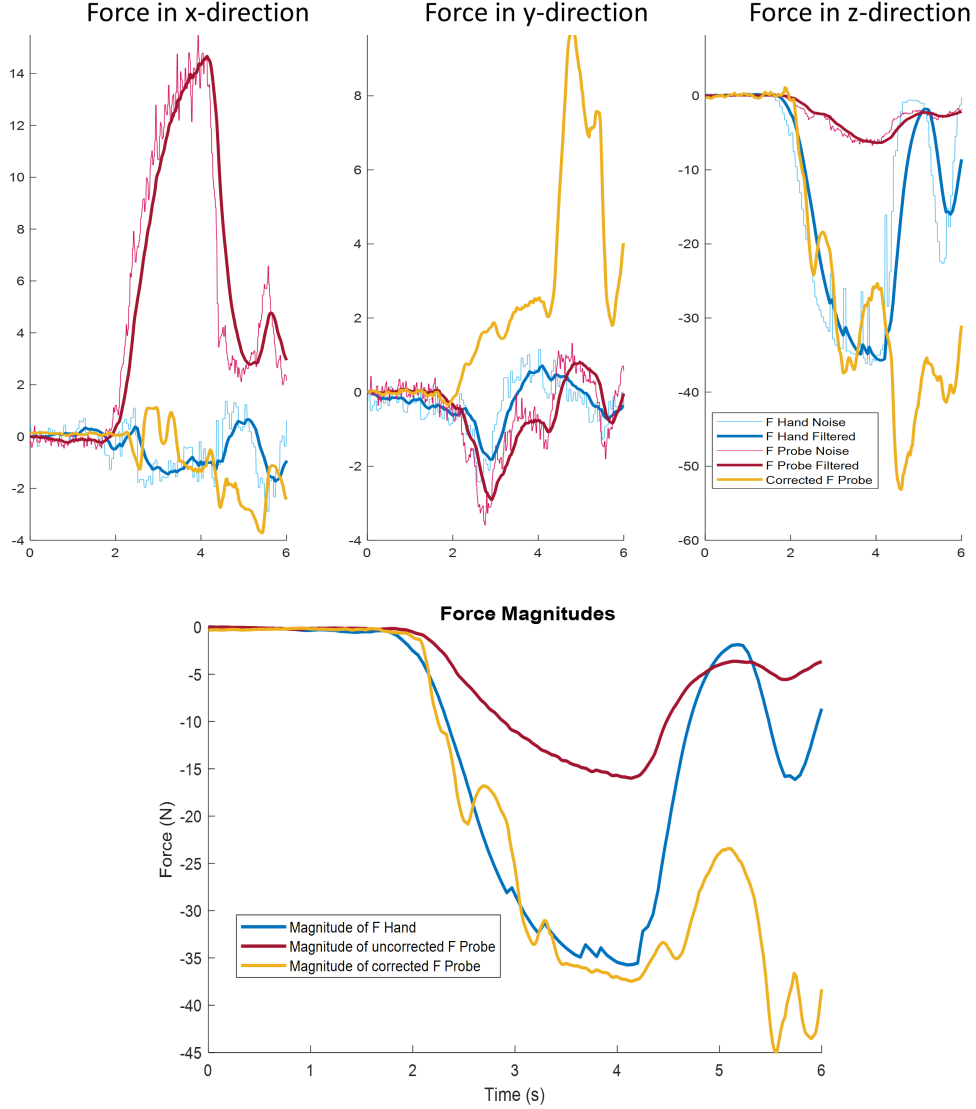


Figure 15: Dataset two of \mathbf{F}_{hand} , \mathbf{F}_{probe} , and the corrected $\hat{\mathbf{F}}_{probe}$ from the neural network, collected on the day after the network was trained.

4.B Gravity Compensation

The gravity compensation procedure described in Section 3.D was performed and resulted in $-mg = -4.7631$, and $p_{COM,xy} = (-0.01250, -0.00033)$.

These parameters were tested on 32 static robot poses with the ideal result being that $0N$ of force and $0N \cdot m$ of torque are expressed after compensation regardless of the uncompensated reading. The results of the compensation are shown below, with green arrows showing significant improvement from compensation and red arrows showing significant worsening from compensation.

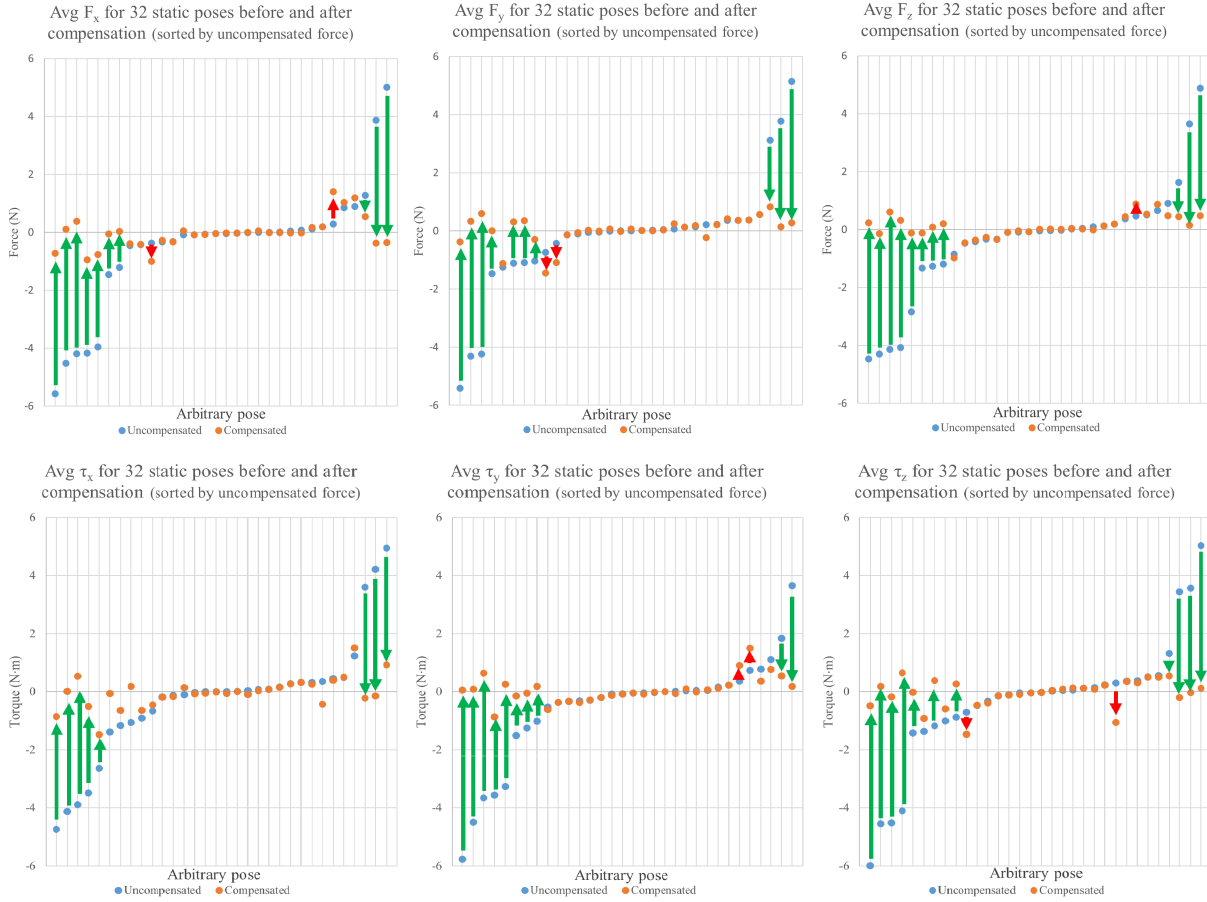


Figure 16: A graphical view of the gravity compensation results. Ideal compensation would result in the blue dots converging to orange dots along the x-axis meaning that no matter the probe orientation, the algorithm could always predict the weight contribution of the tool to the force sensor.

Using data from the same 32 static poses, radar plots were constructed to show which robot orientation caused the most error after compensation.

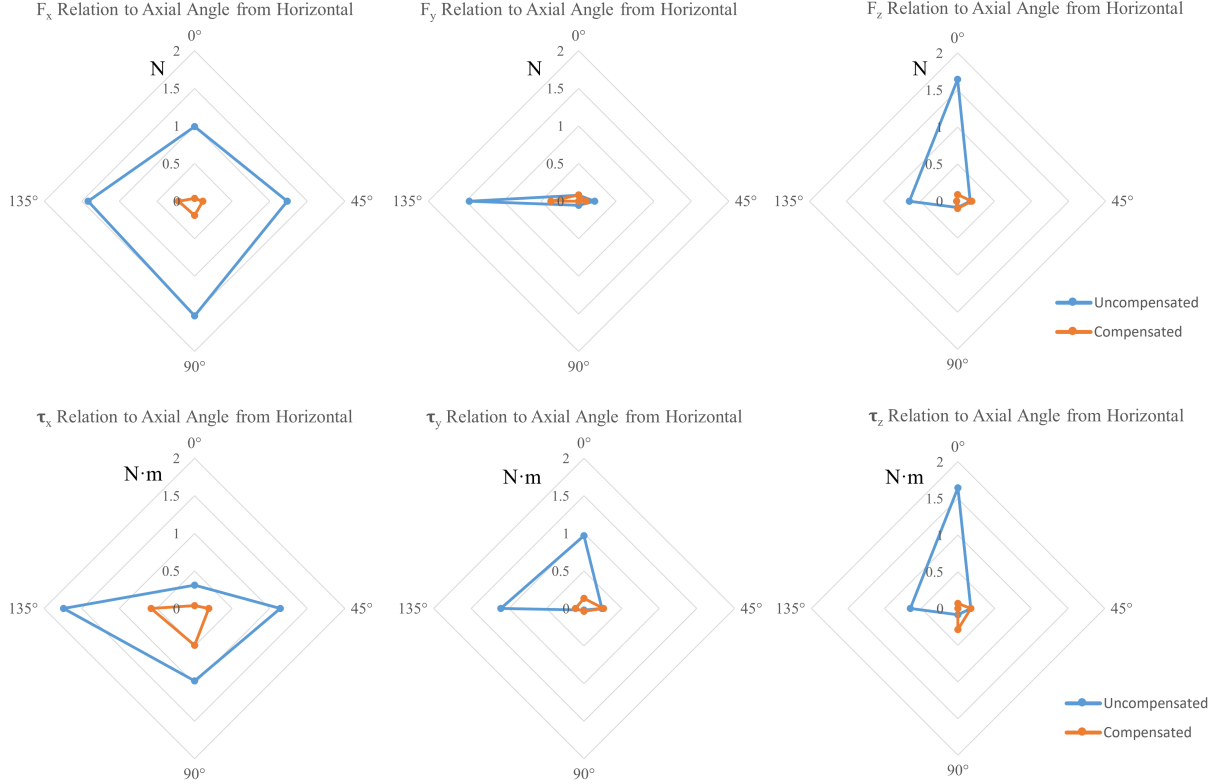


Figure 17: A radar view of the gravity compensation results to show compensation effectiveness for multiple angles. Ideally, these plots would have a uniformly shaped polygon close to the zero force point meaning that compensation was approximately the same for all angles.

4.C Kalman Filtering

Kalman filtering was implemented as described in Section 3.E. Force and torque data was collected for static poses to show noise attenuation, as well as while manipulating the robot around the workspace to show phase-lag. While filtering was applied to \mathbf{F}_{hand} , $\boldsymbol{\tau}_{hand}$, and $\boldsymbol{\tau}_{probe}$, only the results for \mathbf{F}_{hand} will be presented here for brevity since they are most important.

The results of the Kalman filter on \mathbf{F}_{hand} noise while the robot was in a static pose is shown in Fig. 18. The filter reduced the maximum absolute noise for the x-, y-, and z-directions respectively from 1.15 N, 1.06 N, and 0.31 N to 0.58 N, 0.48 N, and 0.12 N. It should be noted that the sensor noise was inherently lower in the z-direction, a factor that led our team to use a different set of admittance control gains and Kalman covariances (with stronger trust in the measured sensor reading) for forces in the z-direction versus those in the x- and y-directions.

Effects of Kalman Filtering on Noise in a Static Pose




Figure 18: The result of the tuned Kalman filter on F_{hand} noise measured while the robot was stationary. The Kalman filter effectively reduced the maximum absolute noise amplitude by a factor of approximately 2.

The Kalman filtering results were also compared to a 5th order Butterworth filter with a cutoff frequency of 5Hz. The results of this comparison are shown below in Fig. 19.




Figure 19: The result of the tuned Kalman filter on F_{hand} noise while the robot was stationary compared to a 5th order Butterworth filter (evaluated in simulation).

Force data was also collected while manipulating the robot in circles around the workspace to measure phase-lag and accuracy. The results of this analysis are shown in Fig. 20. From the plots, it could be graphically determined that there was approximately 10ms of phase-lag added by the filter and an almost negligible amount of undershoot and overshoot. The rectangular shape of the unfiltered force readings can be attributed to the fact that a force packet only arrived every 5ms, and the force value is held until the next packet arrives. The Kalman filter is much smoother since a new value is predicted every 1ms.

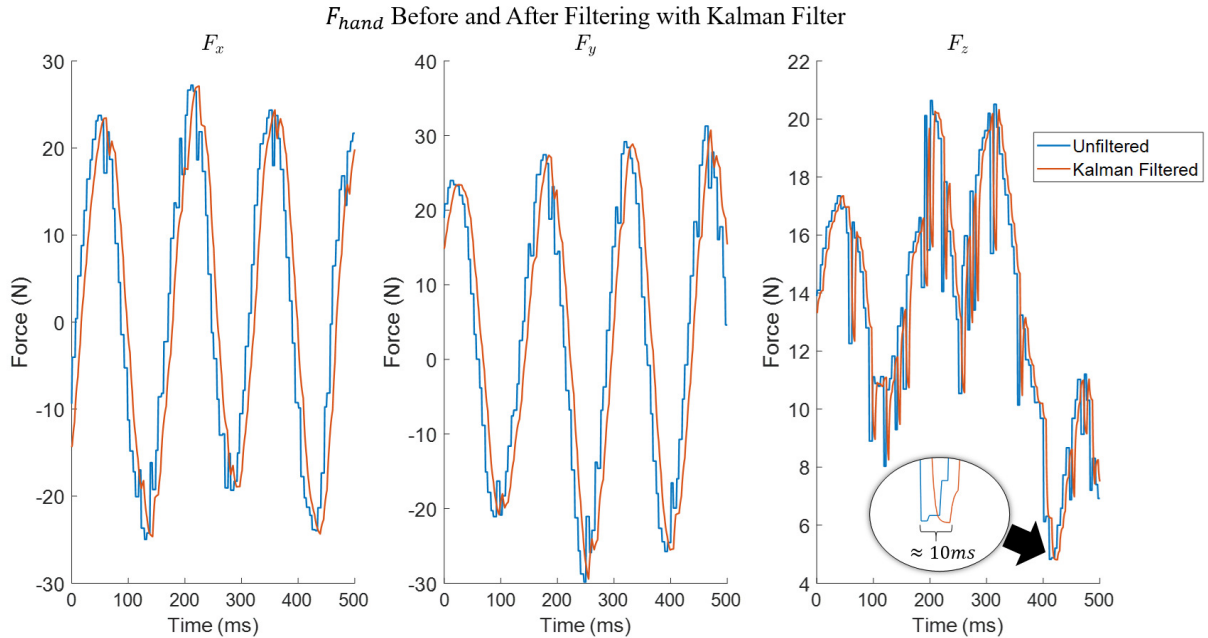


Figure 20: The result of the tuned Kalman filter on F_{hand} while maneuvering the robot around the workspace. The filter introduces approximately 10ms of phase-lag and has very small amounts of overshoot and undershoot.

4.D In-Air Admittance Control

Admittance control was implemented as described in Section 3.F. A plot showing the commanded velocity difference when using the linear gains proposed in this work versus the sigmoidal gains used in [3, 4] is shown below, and uses the same dataset collected to produce the forces shown in Fig. 20. While both curves appear to be smooth, the sigmoidal force-to-velocity conversion clearly undershoots the linear force-to-velocity conversion result. Additionally, the sigmoidal gains cause a nonlinear behavior as the velocity crosses 0 m/s which is where the conversion function enters a deadband zone.

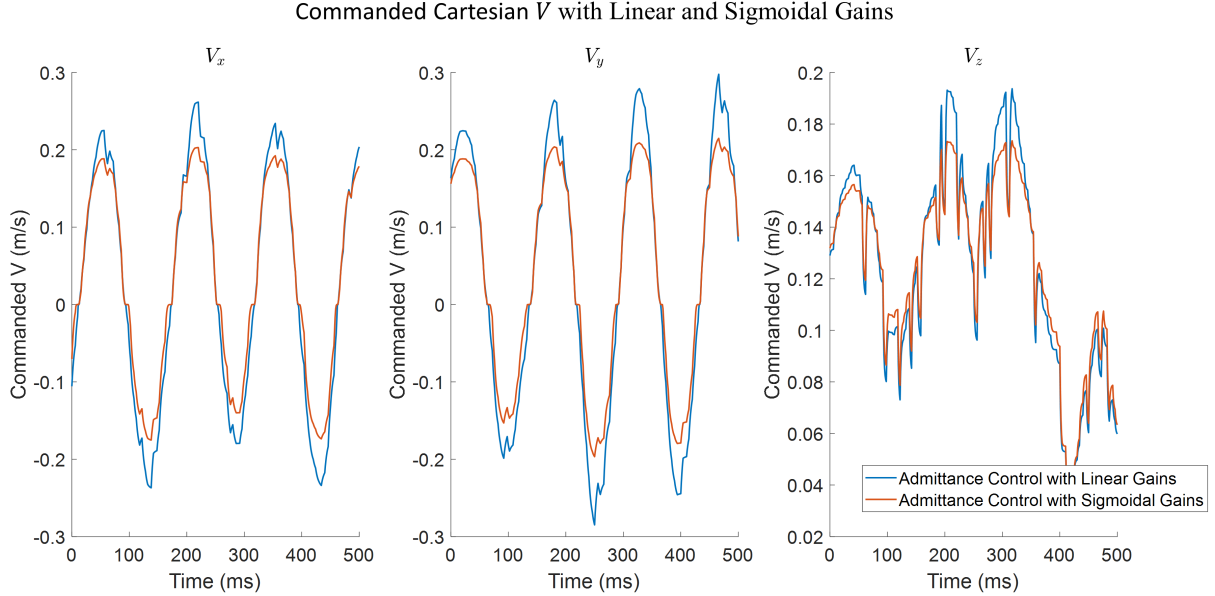


Figure 21: The result of the piecewise linear admittance control algorithm proposed in this work versus the sigmoidal algorithm proposed in [3, 4]. The sigmoidal function undershoots the linear one, and also introduces nonlinear behavior at the 0 m/s crossing.

Unfortunately, despite HIRB approval, a full user study was not able to be conducted due to several hardware hurdles at the end of this project’s implementation. Therefore, there are few quantitative results to describe the “transparency” of admittance control and reduced strain of the user versus manipulating an US probe freehand. It should be mentioned, however, that the admittance control was qualitatively determined to be acceptable by several users who tried the device in an informal “pilot-study.”

5. Discussion

5.A Probe Contact Force Using Neural Networks

The plots in Fig. 13 showing the neural network result on validation and test data suggest that the network characterizes the nonlinearities of the probe forces very well. While a few patterns in the form of “loops” could be spotted on the linear regression plots, the validation and test correlation coefficients were very high especially in the case of the z-direction which is arguably most important when deducing probe forces against a surface.

However, the two plots shown in Figs. 14 and 15 show conflicting test results of the neural network’s ability to correct \mathbf{F}_{probe} into an accurate $\hat{\mathbf{F}}_{probe}$. Fig. 14 shows a near perfect correction of \mathbf{F}_{probe} , whereas Fig. 15 shows an inadequate correction of \mathbf{F}_{probe} in a way that is practically unusable. It was realized that the data in Fig. 14 was collected the same day that the neural network training data was captured, and the data in Fig. 15 was collected the next day. It was hypothesized that the use of the probe by other lab members during the time in-between altered the nonlinear characterization of the probe forces and reducing the effectiveness of the neural network’s predictive power. Since the characterization

of the probe could not persist through perturbation, it would be practically futile to further train the neural network in an effort to wholly capture the probe’s characterization in a generalizable way. Therefore, the neural network approach was abandoned and other ways to determine probe force must be sought.

It was unfortunate to have two force sensors available, but one broken (1-DoF Honeywell) and the other redirecting forces in unpredictable directions (3-DoF Variense). If given more time, refabricating the housing for the 3-DoF Variense sensor in a more rigid material could be useful. Future work should improve upon the housing designs and structural integrity, and possibly use 6-DoF sensing in an effort to better measure probe forces for contact compensation.

5.B Gravity Compensation

The housing parameters solved by the least-squares method matched what was expected. The probe and housing weighed 0.48 kg when placed on a scale which would result in a theoretical $-mg_{theoretical} = 4.7040$, which is very close to the calculated $-mg = -4.7631$. Additionally, the housing is approximately symmetric with the exception of the Variense force sensor located slightly offset in the -x-direction, which accounts for the fact that the calculated $p_{COM,xy} = (-0.01250, -0.00033)$ is slightly offset toward the -x-direction and has negligible offset in the y-direction.

In addition to the parameters seeming logical, the graphical results in Fig. 16 show that adequate compensation was achieved. Almost all of the poses experienced an improvement through compensation and since most cases experienced a compensation to within the noise range of the sensor ($\pm 1N$ or $\pm 1N \cdot m$), the compensation was deemed successful.

The radar plots in Fig. 17 show that the compensation accuracy may be related to pose orientation in some cases. For instance, F_x is approximately equal in error no matter the orientation, whereas F_y is very accurate for all poses except when the y-axis is 135° from the horizontal. This is odd since one would expect the accuracy for 135° and 45° from the horizontal to have the same accuracy. However, it is likely that cable drag plays a role here since the cable protrudes from the housing along the y-axis and was not routed in a way that kept it in a deterministic position. Therefore, the tool weight was felt differently when the y-axis is 135° and 45° from the horizontal. Future work should further consider eliminating cable drag, or making it more deterministic, through a mechanical mechanism.

5.C Kalman Filtering

The plots in Figs. 18 and 19 show that the Kalman filter was successful in attenuating noise for the force sensors while the robot was in a static pose. The noise was approximately cut in half for each direction. This is useful since the deadband zone of the admittance control algorithm (10) can be made narrower since η is smaller versus when using unfiltered data, which makes the motion feel “snappier” and more “transparent” to the user. Additionally, the Kalman filtering appeared to be superior to frequency-domain filtering using a Butterworth filter. Future work could try to use a Butterworth and Kalman filter in series, but this is not guaranteed to produce a better result.

The plots in Fig. 20 show the filtered values while maneuvering the robot in circles. The phase-lag of 10ms seems small, and users did not report feeling it while moving the robot. However, users did notice “inertia” once they stopped applying hand forces to the robot. This could be attributed to the fact that “human-in-the-loop” helps compensate for phase-lag and overshoot while a user is operating the robot, but once the user stops applying forces the lag and inaccuracies become much more noticeable.

5.D In-Air Admittance Control

The in-air admittance control was qualitatively evaluated by mentors and fellow lab members, who provided informal feedback to aid in tweaking the control gains for optimal transparent motion. Typically, while difficult to exactly quantify, users were impressed with the smoothness of the robot motion, especially given the noise characteristics and sample rate of the force sensors being used. Although this project was approved by the HIRB for a user study to collect more definitive quantitative metrics relating to the admittance control performance, the study had to be postponed due to previously-mentioned malfunctioning contact force sensors. The HIRB study planned to have users perform several US scans where stability of probe contact force, stability of acquired US image, muscle exertion via surface electromyography, and survey data would be collected to validate the system performance. Potentially, if a better contact force measurement approach is devised over the Summer, this HIRB-approved study can be conducted in the Fall.

6. Progress Evaluation

In general, the project went according to schedule and without hindrance until late March when the probe force sensor issues began to arise.

6.A Dependencies

Almost all dependencies were met/acquired by the end of February, on schedule, with the exception of acquiring an sEMG system for quantifying the effort exertion difference between using the hand-over-hand control versus the traditional freehand scanning approach. Even so, this dependency was a “nice-to-have” and the user study could have continued with out it.

Dependency	Need	Status	Followup	Contingency Plan	Planned Deadline	Hard Deadline	
Robot	Actuated to provide "power-steering"	Have a working UR5	N/A	If breaks, could seek permission to use Dr. Kazanzides's UR5	2/1	2/1	✓
6DOF F/T Sensor	Admittance control input	Have a working Robotiq FT-150	N/A	If breaks, approach the CS dept to borrow one of their FT sensors	2/1	2/1	✓
Contact Force Sensor	Decouples force from probe on pt.	Have a working 3DoF Variense FSE103	N/A	If broken, continue with load cell or without contact force feedback	2/12	2/28	✓
Ultrasound Probe	Key component for realistic testing	Have a linear probe, several others available in our lab pod	N/A	If disappears, seek permission to use another probe available in lab pod	2/1	4/1	✓
sEMG sensor	Used to measure physical exertion while scanning	Looking to acquire through MUSiC Lab collaborators	Speak with Dr. Boctor	If unable to acquire, testing can still proceed without sEMG data	3/8	4/12	✗
Phantom (non-anatomical)	Something to test the probe on	Acquired	N/A	If disappears, seek permission to use one of the many phantoms present in B08A	2/1	4/1	✓
HIRB First Submission	Testing with subjects	Submitted, trained for human subjects testing, HIPPA	N/A	If not approved in time, we can still perform qualitative validation with sonographers to see if exertion is improved in their expert opinion.	2/22	3/1	✓
HIRB Approval	Testing with subjects	Received brief feedback, crafted an update	N/A	If not approved in time, we can still perform qualitative validation with sonographers to see if exertion is improved in their expert opinion.	3/29	4/12	✓

Figure 22: The final list of dependencies. The yellow boxes indicate where contingency plans were needed, whereas the checks indicate if a dependency was acquired. Some dependencies have a check and a yellow box, indicating that while the item was acquired, it became nonfunctional and a contingency plan had to be enacted.

Although initially fulfilled and given a green checkmark, contingency plans had to be acted upon for several of these dependencies in April as components broke or functioned unexpectedly. These are shaded yellow in Fig. 22. First, although two force sensors were acquired for the purposes of resolving probe forces, both sensors performed very poorly for this application. Due to budgeting and the timing of these failures, the contingency plan was to devise a hack as described in Section 3.F.3 to fake the functionality for the final demo. Second, the team's UR5 robot broke a week before the final poster session. The contingency plan involved asking Dr. Kazanzides to borrow his UR5 robot, which he did graciously donate to the team for a week. While these two contingency plans satisfied the team's short-term needs, a longer-term solution will be needed to continue this work in the future.

6.B Activities and Deliverables

All minimum activities/deliverables, most expected activities/deliverables, and a small bit of the maximum activity/deliverable were completed. An outline of the deliverables and their completion status is shown in Fig. 23, where green indicates fully completed, yellow indicates partially completed or worked on, and red indicates not completed.

	Activity	Deliverable
Min.	C++ interface with robot and dual force sensors to collect data	Datasets for multiple static poses
	Implement rudimentary in-air admittance control, gravity compensation	Video of functionality, graphs showing compensation, code and documentation
Expected	Implement improved admittance control through adaptive Kalman filtering incorporating probe-pt. force feedback	Video of functionality, code and documentation
	Qualitatively and quantitatively evaluate the system with test subjects	Report with graphs and statistical validation
Max.	Virtual fixtures	Video of functionality, code and documentation

Figure 23: A list of all key activities and deliverables for this project, color coded according to completion status. Green means completed, yellow means somewhat completed or worked on, and red means not done.

The team was successful in completing the minimum tasks by the end of February according to the schedule. Most of the expected tasks were completed by early April according to schedule, with two small exceptions. First, the probe-patient contact force compensation algorithm, while incorporated into the programmatic implementation, was never tested due to malfunctioning force sensors. Since the algorithm was programmed and is expected to work in the future with an accurate force sensor, this task has been marked as yellow. Second, quantitative feedback was never collected from test subjects due to the aforementioned malfunctioning force sensors. While qualitative feedback from users was used to tune the control gains and make the motion more transparent, no quantitative metrics were collected from our pilot study and therefore this activity is marked in red. The singular maximum activity of virtual fixtures is highlighted in yellow since a basic virtual fixture was used in Section 3.F.3 to enforce probe-patient contact in the absence of a functional contact force sensor. Although the fixture was a horizontal plane normal to the axis of the probe and assumed a horizontal scanning surface, it is a first step toward the maximum deliverable.

6.C Schedule Adherence

The original schedule for this project was bold but attainable, and was adhered to until late March when unforeseen hardware issues arose. Around the project’s checkpoint presentation, a schedule alteration was made to spend less time on Kalman filtering and more time on contact force sensing. This arose from the fact that the Kalman filtering produced acceptable results faster than anticipated, and the available contact force sensors were outputting strange force values. Without confirmation that both force sensors were practically unusable for this project, especially since they were both used by prior works, the team dedicated

several weeks to developing software solutions that corrected for their unusual values. Unfortunately this resulted in the postponement of a user study past this semester. The eventual, true schedule that resulted is shown in Fig. 24 with documentation and dependency resolution left out for brevity.

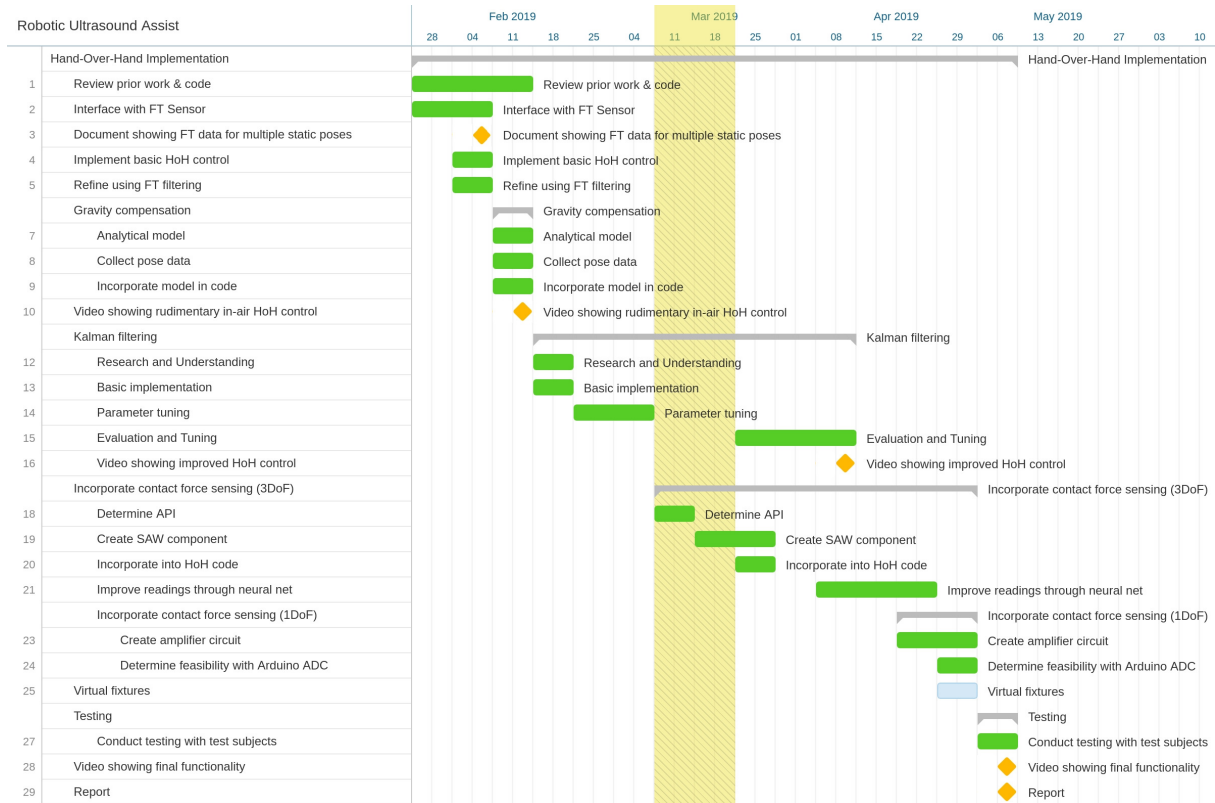


Figure 24: The final weekly schedule for this project, excluding documentation and dependency resolution for brevity. The schedule matches the one proposed at the beginning of the year up until after spring break, where sensor issues caused a delay in the proposed user study and maximum deliverable.

6.D Reflection

Knowing what is known now, it certainly would have been advantageous to identify the issues with the probe force sensors earlier. While the sensors were checked in late January to make sure they were creating readable output, the readings themselves were not verified until late March since improving the transparency of hand-over-hand control was the main focus of this work and the force sensors had been demonstrated as functional in previous papers. Without the hardware issues experienced (or if they were identified earlier), effort would not have been needed in using a neural network to correct the Variense force sensor readings and designing/debugging an amplifier circuit for the Honeywell load cell. This time, regrettably about 4 weeks, could have been spent elsewhere toward a user study, the project's maximum deliverable, and other advanced applications.

With this hindsight behind, the author believes that time was managed well and problems were handled thoroughly and appropriately. The author did everything possible given the

time and scope of materials to work with to produce the maximum deliverable on time and does not regret any scheduling decisions.

7. Conclusion

7.A Significance

The results of this work are significant on two fronts related to the earlier proposed goals. First, the smooth admittance control scheme based around Kalman filtering appears to be promising for reducing muscle strain while scanning and augmenting sonographer scanning ability. While a full user study with sEMG sensing would be needed for complete validation, the preliminary results and qualitative feedback certainly make it appear that the proposed solution would be helpful on these fronts. Second, the functional C++ implementation is important for any future work involving robotics and US at JHU. Since the code was built on the CISST/SAW libraries, this work is very extensible and the general hand-over-hand control framework could be applied to numerous (including non-US) applications that could benefit from hand-over-hand control.

7.B Moving Forward

Before any future work is conducted, the contact force sensor issue must be fixed. Over the few days after this project is due, the team will consult with mentors to determine the best approach. Options being considered include refabricating the existing housing with a less compliant material, redesigning the force sensor housing, and pursuing a new robot with built-in force sensing at the end-effector.

The work performed in this project will then be continued and extended by the author as part of a Master's thesis. The thesis aims to utilize this hand-over-hand US probe control framework in conjunction with a cooperatively controlled transrectal probe to acquire transmission US images for limited-angle tomographic reconstruction to visualize the prostate during cancer screening. The hand-over-hand control is a key enabling technology to align the hand-guided abdominal US probe with the automatically-guided transrectal probe for US transmission imaging. Without robotic precision, such a procedure would be near impossible.

7.C Final Word

Ultimately, this project was a great application of many CIS1 concepts including robot kinematics, solving least-squares problems, gravity compensation, admittance control, determining rotation matrices, and programming with CISST. Applying the classroom knowledge from CIS1 to a real moving project was a beneficial learning experience, and also a lot of fun. While hardware issues inhibited a functional prototype and full user study, the implementation and small pilot study showed promising results that will be further verified in future work.

References

- [1] A. Schoenfeld, J. Goverman, D. Weiss and I. Meizner, “Transducer user syndrome: an occupational hazard of the ultrasonographer”, *European Journal of Ultrasound*, vol. 10, no. 1, pp. 41-45, 1999. Available: 10.1016/s0929-8266(99)00031-2.
- [2] T. Rousseau, N. Mottet, G. Mace, C. Franceschini and P. Sagot, “Practice Guidelines for Prevention of Musculoskeletal Disorders in Obstetric Sonography”, *Journal of Ultrasound in Medicine*, vol. 32, no. 1, pp. 157-164, 2013. Available: 10.7863/jum.2013.32.1.157.
- [3] R. Finocchi, “Co-robotic ultrasound imaging: a cooperative force control approach”, The Johns Hopkins University, 2016.
- [4] R. Finocchi, F. Aalamifar, T. Fang, R. Taylor and E. Boctor, “Co-robotic ultrasound imaging: a cooperative force control approach”, *Medical Imaging 2017: Image-Guided Procedures, Robotic Interventions, and Modeling*, 2017. Available: 10.1117/12.2255271.
- [5] H. K. Zhang, R. Finocchi, K. Apkarian and E. M. Boctor, “Co-robotic synthetic tracked aperture ultrasound imaging with cross-correlation based dynamic error compensation and virtual fixture control,” *2016 IEEE International Ultrasonics Symposium (IUS)*, Tours, 2016, pp. 1-4. Available: 10.1109/ULTSYM.2016.7728522
- [6] T. Fang, H. Zhang, R. Finocchi, R. Taylor and E. Boctor, “Force-assisted ultrasound imaging system through dual force sensing and admittance robot control”, *International Journal of Computer Assisted Radiology and Surgery*, vol. 12, no. 6, pp. 983-991, 2017. Available: 10.1007/s11548-017-1566-9.
- [7] F. Dimeas and N. Aspragathos, “Learning optimal variable admittance control for rotational motion in human-robot co-manipulation”, *IFAC-PapersOnLine*, vol. 48, no. 19, pp. 124-129, 2015. Available: 10.1016/j.ifacol.2015.12.021.