

Low-Profile 6-Axis Differential Magnetic Force/Torque Sensing

David Black, Nicholas Ranga Pradnyawira, Mika Nogami, Amir Hossein Hadi Hosseinabadi, Septimiu Salcudean

Abstract—Force/torque sensing on hand-held tools enables control of applied forces, which is often essential in both tele-robotics and remote guidance of people. However, existing force sensors are either bulky, complex, or have insufficient load rating. This paper presents a novel force sensing modality based on differential magnetic field readings in a collection of sensor modules placed around a tool or device. The instrumentation is easy to install and low profile, but nonetheless achieves good performance. An optimization-based design procedure is also introduced. The modeling, simulation, and optimization of the force sensor are described and then used in the electrical and mechanical design and integration of the sensor into an ultrasound probe. Through a neural network-based nonlinear calibration, the sensor achieves average root-mean-square errors of 0.60 N and 0.037 Nm compared to an off-the-shelf ATI Nano25 sensor, which are 1.17% and 1.63% of the full-scale range respectively. The sensor has an average noise power spectral density of less than 0.0001 N/ $\sqrt{\text{Hz}}$, and a 95% confidence interval resolution of 0.0086 N and 0.0628 Nm. The practical readout rate is 1.3 kHz over USB serial and it can also operate over Bluetooth or Wi-Fi. This sensor can enable instrumentation of manual tools to improve the performance and transparency of teleoperated or autonomous systems.

Index Terms—Force Torque Sensing, Hall effect, Teleoperation, Haptics, Calibration, Ultrasound

I. INTRODUCTION

Force/torque sensing instrumentation of manual tools has many useful applications. For example, it can be valuable in medical ultrasound (US) where force is an important parameter to control, as it determines image quality and the visibility of anatomical structures. For teleoperation and remote guidance, haptic feedback is essential to make the interaction intuitive and realistic for the operator. Force/torque measurement is therefore particularly important in tele-robotic US procedures [1]–[5], or for “teleoperating” novice people in performing tests through a mixed reality interface [6], [7]. In addition, sonographers often face increased incidence of musculoskeletal injury due to the awkward application of loads during exams [8]. This can lead to discomfort for the patients as well, both of which could be reduced through safety monitoring with a force sensor. Furthermore, for US training [9], [10], skill assessment [11], and tele-guidance [12], [13], the applied forces are valuable. Autonomous robotic US also uses force control [14], [15], and for learning from demonstration-based approaches to AI-guided US, data including forces must be collected from US exams [16], [17].

However, force/torque measurement on an US probe is difficult as radiologists resist added bulk, weight, and cable pull. Moreover, the sensor cannot be placed on the face of the probe, in the direct load path, as it would disrupt image

formation. The usual approach for instrumenting US probes has been placing off-the-shelf (OTS) force/torque sensors between the probe and an external shell which is held by the user [2], [18]–[21]. This approach makes the probe bulky, heavy, and difficult to grasp, is expensive, and introduces cable pull. Because the force sensor cannot be placed near the US transducer’s face which is in contact with the patient, it requires high torque capability, leading to high cost and limited availability. A different approach by Huang et al. placed small piezoresistive pressure transducers on either side of an US transducer array [22], but this measures only one or two degrees of freedom (DOFs) of force information and may interfere with the imaging. Most of the robotic US systems mentioned in the last paragraph have a force sensor on the robot end-effector, which is of course infeasible for manual US. It is also limiting if a generic grasper is used to interact with different objects. The same factors are true for any other manual tool, where the tip is used to interact with the environment and should not be covered, and the handle should not be made bulky or heavy.

We thus undertook to design a low-cost, low-profile, easy-to-fabricate and use sensing solution for manual tools to enable force feedback and control in remote guidance or robotic teleoperation, as well as the many other applications. There are multiple potential modalities for force/torque sensing, reviewed in [23], [24]. These include strain gauges [25], fibre Bragg gratings [26], elastomeric transducers [27], piezoresistive pads [28], optical deflection sensing [29], and capacitive sensing [23]. Each modality was evaluated and/or tested [30], but each had issues. In particular, optical sensing requires precise fabrication to align the slit with the LED and photodiode, which involves a relatively complex adjustment step, and the components use more space than is desirable around a manual tool [31]. Piezoresistive transducers require precise pre-loading so all axes contact properly at all times (i.e. they cannot measure negative pressures), and effective elastomeric sensors are not yet commercially available [32]. We tested several models of piezoresistive transducers which showed large hysteresis and inter-axis coupling [30]. Conversely, strain gauges require careful surface preparation, material selection, and manufacture of the flexure, leading to higher cost and complexity [25]. Additionally, installation and removal would be time consuming, which could impede the sanitization of medical tools, eg. US probes, between procedures. Capacitive sensing again involves more complex electronics and careful fabrication because high resolution position sensing is normally only possible with small distances between the capacitor’s plates [23].

Magnetic sensors, including Hall effect and anisotropic magneto-resistance (AMR) sensors, have also seen limited use for force measurement. Jones et al. and Kyberd et al. position a Hall effect sensor near a magnet on a flexible structure such as the fingertip of a robotic grasper to sense when a single-axis force is applied [33], [34]. Initial work on a similar approach using a 3-axis AMR sensor is introduced by Yu et al. [35]. AMRs are also used for position sensing [36], and present some advantages over Hall effect sensors, including better sensitivity. However, they require a Wheatstone bridge configuration which would add significantly to the size and complexity of a sensor. An interesting human finger tip force sensor using the Hall effect is introduced in [37], but it is specifically usable only with a human finger. Conversely, a 3-axis force sensor with Hall effect sensors and magnets was designed in [38]. Again, this uses single Hall effect sensors axially offset from magnets. To achieve sufficient signal, the sensor relies on large deflections and cannot be miniaturized. As a result, the presented device is bulky and relatively heavy, with a small full-scale range, and the response is non-linear.

Instead, we previously introduced a multi-axis force sensing approach using Hall effect sensors in a differential configuration [30]. In that preliminary work, the sensing modality was described and shown through simulation, and in one-, and three-axis tests, to be effective for force and torque sensing. The best-case deflection resolution of a single transducer was found to be 856 nm, and for a 3-axis jig with a linear calibration, the RMS error was approximately 10.37%.

In this paper, we build on the prior work through the following contributions. We introduce

- A physical model, simulation, and automated optimization-based approach for the force sensor integration into an arbitrary tool,
- The electrical design, communication architecture, and firmware of miniaturized printed circuit boards (PCBs) for the sensor,
- A novel 3D-printed suspension design, characterization, and subsequent selection using the simulation,
- The mechanical design, fabrication, and integration of the force sensor into an ultrasound probe, informed by the optimization,
- Linear and neural network-based calibration of the ultrasound probe force sensor,
- Analysis of the calibration, speed, noise characteristics, sensor orientations, and mechanical design, and verification of the simulation and optimization.

In the following sections, we first briefly explain the modality and sensing concept (Section II-A) before describing the above contributions, designs, and tests in Sections II, III, and IV respectively. The final design and integration into an ultrasound probe is shown in Fig. 1.

II. METHODS

A. Differential Magnetic Force Sensing

As previously described [30], differential magnetic force sensing (DMFS) uses a collection of n sensor modules placed around a tool or device. Each sensor module consists of two



Fig. 1. Design and integration of 6-DOF force sensor onto a C3HD 3 ultrasound probe (Clarius, Vancouver, BC), with inner scaffold, compliant suspension, outer shell, and electrical integration.

adjacent Hall effect sensors, mounted on a thin shell around the tool, and a small permanent magnet mounted opposite the sensors on the tool itself. The shell is attached to the tool by a compliant suspension. When no load is present, the magnet of each sensor module is centered between its two Hall effect sensors. However, when the user applies a force to the shell, the tool deflects relative to the shell, and thus the magnets move relative to their sensor pairs, which causes a voltage response. An overview is shown in Fig. 2. By calibrating the sensor module deflection outputs with the stiffness of the suspension, forces and torques can be computed.

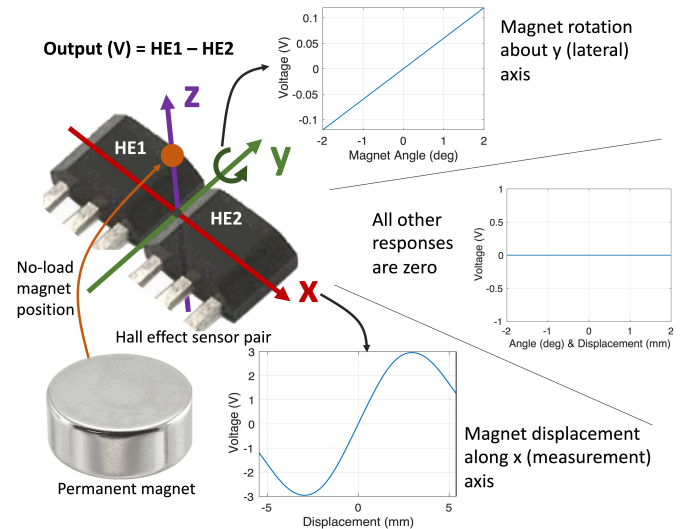


Fig. 2. Measurement concept of DMFS sensors. Taking the difference of the Hall effect sensor outputs, all displacements cancel out except along the x axis, where it is sensitive and very linear in the ± 2 mm range ($R^2 = 0.999$ [30]). For rotations about y , the response is small and the magnet is highly constrained. The magnet is offset 1-1.5 mm along the z (axial) direction. Motion along that axis can be detected by summing the sensor outputs.

In particular, taking the difference between the two Hall effect sensor outputs, the response is highly linear for magnet displacements across the sensors (i.e. the x or measurement

axis in Fig. 2). For all other axes, the response cancels out or is very small. Taking the sum of the two sensor outputs gives information about the axial (z axis) offset, but also the lateral (y -axis). Hence, each sensor module is considered a 1-DOF transducer, though there is some information about additional DOFs. Thus, approximately n modules are needed for n -axis force/torque measurement.

B. Physical Model

Before designing the force sensors, the concept was modeled in detail, and the model was verified through simple tests with one Hall effect sensor, followed by single-SM tests and finally a 3-DOF test [30]. This modeling is now extended to a general multi-DOF case to aid in the design of the 6-DOF configuration.

To measure the forces and torques applied to an arbitrary manual tool, we place a thin shell around the handle of the tool, where the user grips it. The shell is attached to and separated from the tool using a compliant suspension consisting of m discrete connection points. For example, a connection point could be a rubber pad. Between the tool and the shell are also n SMs.

The suspension members can be placed in any reasonable configuration that supports the shell relatively uniformly and symmetrically. However, the sensor locations affect the achievable accuracy and numerical stability of the calibration. Thus, we must intelligently select their positions. To do so, we first develop a mathematical model of the forces and displacements in the suspension when the shell is held fixed and a force is applied to the tool. We then develop a simulation in MATLAB that implements this model, taking in a set of sensor positions and applied forces, performing a calibration, and outputting the mean-squared error (MSE) of the simulated ‘‘measured forces’’. This simulation can be run repeatedly for different sensor positions, and the errors compared to determine the optimal sensor configuration.

1) *System Mathematical Model:* In the following, we use coordinate-free notation for simplicity. Practical considerations of coordinate frames for the simulation are described in Appendix I. Under an applied wrench $\mathbf{w}_j = [\mathbf{f}_j^\top \ \boldsymbol{\tau}_j^\top]^\top$ (force \mathbf{f}_j and torque $\boldsymbol{\tau}_j$), the j^{th} suspension member undergoes a 6-axis deformation, $\boldsymbol{\delta}_j = [\boldsymbol{\delta}_t^\top \ \boldsymbol{\delta}_\theta^\top]^\top$, according to Hooke’s law (for small deformations):

$$\mathbf{w}_j = \mathbf{K}_j \boldsymbol{\delta}_j \quad (1)$$

Where \mathbf{K}_j is the rank 2 compliance tensor of the sensor element. The matrix value of \mathbf{K} in a practical coordinate system is described in Appendix I, and we performed tests to find its elements, described in Section III-B1. The rotational part of the deflection, $\boldsymbol{\delta}_\theta$, is a vector of small angles since the motion is very constrained. Additionally, we assume the shell and tool are rigid bodies, so we can express the individual deflections in terms of the overall deflection at the tool tip:

$$\boldsymbol{\delta}_j = \begin{bmatrix} \boldsymbol{\delta}_t + [\boldsymbol{\delta}_\theta]_\times \mathbf{x}_j \\ \boldsymbol{\delta}_\theta \end{bmatrix} \quad (2)$$

where \mathbf{x}_j is the position of the j^{th} suspension element, and $[\mathbf{v}]_\times$ is the skew symmetric matrix equivalent to the cross

product such that $[\mathbf{v}]_\times \mathbf{u} = \mathbf{v} \times \mathbf{u}$. Now suppose a wrench, \mathbf{W} , is applied at the tool tip. From statics:

$$\mathbf{W} = \sum_{j=1}^m \begin{bmatrix} I_3 & 0_3 \\ [\mathbf{x}_j]_\times & I_3 \end{bmatrix} \mathbf{w}_j \quad (3)$$

where I_k and 0_k are the $k \times k$ identity matrix and zero matrix respectively. Substituting in Equations 1 and 2, we find

$$\mathbf{W} = \sum_{j=1}^m \begin{bmatrix} I_3 & 0_3 \\ [\mathbf{x}_j]_\times & I_3 \end{bmatrix} \mathbf{K}_j \begin{bmatrix} \boldsymbol{\delta}_t + [\boldsymbol{\delta}_\theta]_\times \mathbf{x}_j \\ \boldsymbol{\delta}_\theta \end{bmatrix} \quad (4)$$

Equation 4 can be solved numerically for $\boldsymbol{\delta}_t$ and $\boldsymbol{\delta}_\theta$, taking the coordinate frames into account (Appendix I). Having solved this, the result can be substituted into Equation 2 to determine the $\boldsymbol{\delta}_j$ ’s. These give the suspension member wrenches, \mathbf{w}_j by Equation 1.

In the force sensor, the sensor and suspension element locations are not required to coincide. The sensor locations, $\{\mathbf{s}_i\}$, can be arbitrary, though constrained by practical mounting considerations, and the displacement at the sensors can be determined by substituting \mathbf{s}_i into Equation 2 instead of \mathbf{x}_j . Furthermore, displacement sensor elements generally measure along just one axis in addition to the normal direction. Thus, one must define the sensitive axes, $\hat{\mathbf{s}}_i$, in addition to the sensor positions. The normal directions, $\hat{\mathbf{n}}_i$, are given by the sensor positions and the tool geometry. To simulate a measurement, the calculated $\boldsymbol{\delta}_i$ must be projected onto the normal and sensitive axes:

$$\begin{bmatrix} \delta_s \\ \delta_n \end{bmatrix} = \begin{bmatrix} \boldsymbol{\delta}_i \cdot \hat{\mathbf{s}}_i \\ \boldsymbol{\delta}_i \cdot \hat{\mathbf{n}}_i \end{bmatrix} \quad (5)$$

To make this more realistic, however, we used data from the single SM test [30] to create a look up table (LUT) of SM output as a function of magnet displacement, $\boldsymbol{\delta}_j$. Readings from both Hall effect sensors were recorded for displacement sweeps across the Hall effect sensor pair at axial offsets of 0.5 to 2.5 mm, and lateral offsets of 0.0 to 1.5 mm from centre, both in 0.5 mm increments. Across the sensor pair, the displacement was measured in 0.1 mm increments. A cubic spline interpolation between the measured points was used to compute the sensor responses at a given point. This is illustrated in Fig. 9.

2) *Force Measurement Simulation:* It is now possible to formulate a simulation algorithm with the model. Suppose we have m suspension members and n sensing elements, and a three dimensional mesh of the tool. This is essential in any case to design the shell, and can be created in computer aided design (CAD) software or by using a 3D scanner.

First we choose the desired suspension positions on the mesh, and thus also obtain their orientation. A constraint set, \mathcal{D} is defined for the possible SM positions. The constraint set is application dependent. In the example of an ultrasound probe as seen in Fig. 9, one could constrain the configuration to have two sensors per wide face and one on either side, and to have a certain minimum spacing between sensors. Once the SM locations are randomly initialized within \mathcal{D} , the program can solve the equations presented in Section II-B1 in order to simulate an applied force and subsequent measurement. This is shown in Algorithm 1.

Algorithm 1 Calibration Simulation

Generate a random matrix of N wrenches,
 $W = [\mathbf{w}_1, \dots, \mathbf{w}_N]$
for $\mathbf{w}_k \in W$ **do**
 Solve Equation 4 for δ_t
for $i \in \{1, \dots, n\}$ **do**
 Substitute δ_t and \mathbf{s}_i into Equation 2 to find δ_i
 Use the LUT to simulate measurement
end for
 Place the measurements in column k of $X \in \mathbb{R}^{2n \times N}$
end for
 Determine mapping, C , from X to W (calibration):
 $C = WX^\top(XX^\top)^{-1}$
return RMS error of calibration and $\text{cond}(C)$.

The calibration in the second last step of Algorithm 1 can be achieved in a myriad of ways. The simplest, which assumes the system is linear, is to determine a calibration matrix, C , as the least squares linear mapping between the set of SM measurements X and the actual measured forces W :

$$\begin{aligned} W &\approx CX \\ C &= WX^\top(XX^\top)^{-1} \end{aligned} \quad (6)$$

Alternatively, a relatively shallow neural network can be trained to learn the nonlinear mapping, leading to far more accurate results, as shown in [39]. This should be done in the final design of the instrumented tool, but it is impractical to re-train a neural network repeatedly during the sensor position optimization step. Hence, we use the linear approach during the optimization.

3) *Sensor Position Optimization*: With the simulation in place, it is possible to optimize the locations of the sensors to obtain the most accurate and numerically stable calibration. The simulation is run repeatedly for different sensor configurations, S , and measurement axes, $\{\hat{\mathbf{s}}_i\}$. The calibration matrix, $C(S, \{\hat{\mathbf{s}}_i\})$, is determined for each trial, and the calibration error and condition number of C are evaluated for each.

Bicchi describes the importance of the condition number [40], showing that the *a priori* estimate of the relative measurement error, e_p , is given by

$$e_p \leq \lambda \frac{\kappa(C)}{1 - \epsilon \kappa(C)} \quad (7)$$

where ϵ is very small and λ is a constant that depends on the transducers. Thus, a large condition number can lead to poor accuracy even with careful calibration [41]. As a rule of thumb, k digits of accuracy are lost when $\kappa(C)$ is 10^k [42].

The configuration with the lowest mean-squared error (MSE) and condition number is thus chosen according to the cost function in Equation 8. Here \mathbf{x}_k is the k^{th} column of X , i.e. the k^{th} measurement, $\kappa(C)$ is the condition number, and α, β are weights.

$$\begin{aligned} w(S, \{\hat{\mathbf{s}}_i\}) = \alpha \frac{1}{N} \sum_{k=1}^N \|\mathbf{C}\mathbf{x}_k - \mathbf{f}_k\|_2^2 \\ + \beta \log_{10}(\kappa(C)) \end{aligned} \quad (8)$$

A reasonable worst-case calibration MSE is if $C = 0$; then the MSE becomes the mean squared force magnitude. In Fig. 13, for example, this is 237. The condition number, on the other hand, may vary by several orders of magnitude, so it is reduced to the scale of the MSE using the logarithm. Though the condition number is important, as described above, it only works for non-redundant designs - i.e. n -axis force sensors with exactly n transducers [40]. As the simplest, smallest design is desirable in our case, we do not wish to introduce extra sensor modules, so this is acceptable. However, the assumption of linearity is strong and later sections compare the performance to a non-linear, neural network-based calibration. Hence, we also include the calibration error itself in the cost function.

Thus we arrive at the optimization problem in Equation 9.

$$\begin{aligned} S^*, \{\hat{\mathbf{s}}_i^*\} = \min_{S, \{\hat{\mathbf{s}}_i\}} w(S, \{\hat{\mathbf{s}}_i\}) \\ \text{s.t. } S \in \mathcal{D} \end{aligned} \quad (9)$$

This is not possible to solve analytically because there is no analytical expression for $X(S, \{\hat{\mathbf{s}}_i\})$ or $C(S, \{\hat{\mathbf{s}}_i\})$, which result from repeatedly solving a large set of equations. Thus, we must rely on a search algorithm to explore the solution space.

Though each sensor moves on a relatively small two dimensional surface, there are n SMs, so the algorithm is searching in up to $2n$ -dimensional space. Thus, a brute force method is impractical. Further, the optimal sensing axis, $\hat{\mathbf{s}}_j$, of each transducer is unknown, which adds to the complexity. Depending on the application, it makes sense to strongly constrain the measurement axes. We constrained the sensing axes to be just horizontal or vertical, tangent to the probe at the given point.

The optimization problem was solved using simulated annealing [43]. This is described in Appendix II.

III. ELECTROMECHANICAL DESIGN AND SOFTWARE




Fig. 3. Sensor module (top) and master (bottom) PCBs. The low-pass filter circuit for one Hall effect sensor on the sensor module is also shown.

A. Electrical Design

Two PCBs were designed for the force sensors: a sensor module (SM) board and a master board. A 6-axis force/torque sensor uses a minimum of 6 SMs and one master. These are described below, and the PCBs are shown in Fig. 3.

Master PCB

The master board takes in all the readings from the SMs, combines and processes them on an ESP32-WROOM-32e microcontroller, and sends the resultant force/torque measurements to the host PC or other device. An inertial measurement unit (IMU) is also included on the board for pose estimation [44], and communicates over I2C with the ESP32. The microcontroller has two cores, allowing parallel processing, and two SPI drivers. Each SPI driver handles three of the SMs, thus doubling the measurement rate compared to one SPI driver handling 6 SMs in series. The SPI communication includes data-in and data-out (MISO and MOSI in Fig. 4 - Master In/Out, Slave Out/In) buses, as well as serial clock (SCLK), power (+5V), and ground (GND). These are shared among the SMs in a multidrop configuration while a separate chip select (CS) line goes to each individually. In this way, the wiring can be greatly simplified (See Section V). The layout is shown in Fig. 4.

A universal asynchronous receiver/transmitter (UART) port is available to connect the sensor to a PC using a UART-USB adapter. However, the ESP32 additionally has Wi-Fi as well as Bluetooth classic and low energy (BLE) capabilities. We wrote firmware to communicate the sensor readings to a HoloLens 2 (Microsoft, Redmond, WA) over BLE for a teleoperation application [6], [45] but found that the rate was too low. Instead, UART was used for all the presented tests. Other options are discussed in Section V. The configuration of the master PCB and its communications is shown in Fig. 4.

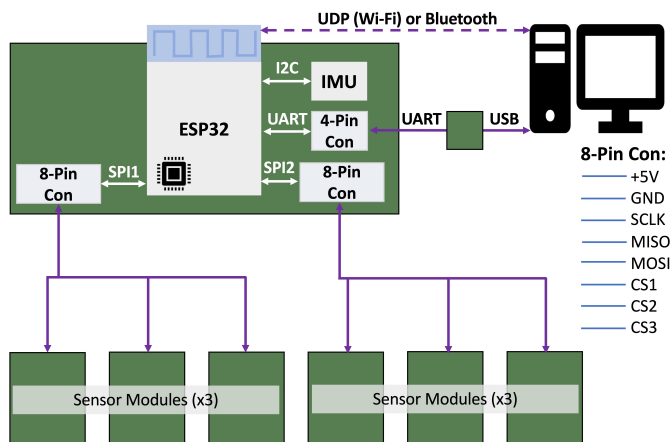


Fig. 4. Communication architecture of sensor modules and master PCB. The pins of the 8-pin SPI connectors are shown on the right.

Sensor Module PCBs

Each SM consists of two Hall effect sensors, used for the differential measurement, a low-pass filter for each sensor, a 12 bit, 100ksp/s analog-to-digital converter (ADC), and a connector for communication with the master PCB. The

filtering and the ADC were placed in close proximity to the sensor to reduce noise.

A passive second order low-pass filter with the component values shown in Fig. 3 was chosen after SPICE simulation with various passive and active configurations due to its simplicity and effectiveness. The -3 dB cut-off frequency was set to 400 Hz since interactions with stiff environments can lead to high force bandwidths. The phase delay is $< 1^\circ$ until approximately 10 Hz, which is the maximum typical for forces applied by the human hand [46].

The SMs are designed as squares with the Hall effect sensors centrally located so that the SM can be rotated by multiples of 90° to change the measurement axis without changing the mounting area.




Fig. 5. Mechanical components of ultrasound force sensor, including inner scaffold and outer shell (A), suspension elements (eg. D), and the complete assembly (B). C shows the UART port for communication with a PC.

B. Mechanical Design

Given the sensor electronics and optimal SM locations, we designed the mechanical components. These were subject to several design objectives, listed below.

- Ease of assembly - the sensor should be easy and fast to install and remove from the object. For ultrasound, this allows disinfection between procedures.
- Ergonomics - the final device should be as low-profile as possible, adding minimal bulk to the ultrasound probe.
- Force range - the probe should displace about 1-1.5 mm (the linear range of the sensor) for a 25 N applied force, as is a typical maximum in ultrasound procedures [30].
- Stiffness - for the sensing concept to work, the outer shell must deflect much less than the deflections in the SMs.

To this end, the ultrasound probe force sensor consists of an outer shell containing the SMs and master PCB, and an inner scaffold containing the magnets. The mechanical parts are shown in Figs. 5 and 6. The inner scaffold clamps directly onto the ultrasound probe while the outer shell connects to the scaffold through the suspension. All three were fabricated




Fig. 6. Ultrasound force sensor mechanical and electrical integration and mounting on calibration jig. The SMs are labeled numerically for later reference. The master PCB shown here has a separate port for each SM, and each was wired separately. Significant improvements in the wiring and PCBs are discussed in Section V.

with a Form 3+ stereolithography 3D printer (FormLabs, Somerville, MA), using different materials.

The scaffold, as it is completely supported by the transducer and has no structural role, was printed from 2 mm thick FormLabs standard white resin. The shell was printed from FormLabs Rigid 1000 polyurethane resin for toughness and rigidity, with a tensile modulus of 920 MPa and bending modulus of 750 MPa. The main wall thickness was 2.5-3 mm, with some sections thicker or thinner as needed. This feels very solid and we did not observe any flexing. The rigidity of the shell was tested using finite element analysis (FEA) in SolidWorks (SolidWorks Corp., Waltham, MA), by placing elastic supports with the known suspension stiffness values, and applying a 25 N force on one side of the shell. The results are shown in Section IV-F.

Finally, the suspension was also 3D printed from FormLabs Elastic 50A resin, with 50A Shore durometer. The design of the suspension is described in Section III-B1. The shell and scaffold were based on a CAD model of the ultrasound probe, and prototyped several times to ensure precise fit. Both shell and scaffold were fabricated in two halves. The scaffold halves snap together by friction fit tabs and are held by the suspension elements. Small magnets embedded in the edges of the outer shell hold the two halves in place (see Fig. 6), though the friction from the suspension flanges and the grip of the user make this relatively superfluous. Throughout testing, the design never opened or shifted unintentionally. It was found to be easy and quick to install, comfortable to hold, robust, and stiff.

1) *Suspension Tests*: Printing the suspension allowed us to integrate it easily into the design by fabricating the suspension element and mounting flanges as one contiguous part. Furthermore, the stiffness could be tuned to the desired level




Fig. 7. Suspension test jig and sample 3D-printed elements showing the inner and outer diameter.

by modifying the part's shape. In particular, a hollow cylinder shape was used with configurable inner diameter (ID) and outer diameter (OD). This is shown in Fig. 7. A flange with two mounting holes was included on either end of the cylinder so it could be screwed into the scaffold and shell. The flanges were offset by 90° to facilitate installation with a screw driver.

To determine the stiffness matrix, K , described in Section II-B, a test jig was created for the suspensions. The measured K value could then be input to the simulation to determine if the deflection response was as desired, and thus a suspension element shape could be selected. The test jig is shown in Fig. 7 and consisted of a manual 3-axis stage on which an ATI Nano25 force sensor (ATI Industrial Automation, Rochester Hills, MI) was mounted, and a fixed flange on which a mounting plate was attached. The suspension element was connected to the mounting plate and an adapter on the ATI force sensor.




Fig. 8. Test jig for torsional rigidity of suspension elements.

The stage was then moved by incremental, known offsets, and the resultant force was measured. A similar setup with a rotational stage was used to measure the rotational stiffness, as shown in Fig. 8. The results are shown in Section IV-A.

C. Integration

The integration of the electrical and mechanical parts on a 3D-printed dummy ultrasound probe as part of the calibration jig is shown in Fig. 6. As explained in Section III-A, the wiring has been cleaned up significantly since these tests, so separate cables are not needed for every SM (See Fig. 16).

The installation process on the ultrasound probe is simple and fast, taking approximately 1-2 minutes. First, the inner scaffold (with magnets installed and suspension elements screwed to one half) is snapped onto the probe, and the suspension elements are screwed to the other half. This takes the majority of the time. Next, the probe with the scaffold is placed in one half of the outer shell, and the other half is fitted on top, making sure the suspension elements are sitting properly. This takes just a few seconds.

The master PCB can then be plugged in to any USB port to receive power. A reset button on the back of the PCB, accessible through a small hole in the outer shell, is pressed to start the firmware. In case of an issue with the firmware, the reset button is always accessible, and the PCB can be flashed with new firmware via the USB cable.

The firmware, written in C using the Espressif IDF and ESP libraries, runs a real-time operating system (FreeRTOS) on the dual-core CPU. This enables multi-threaded parallel processing of the sensor readings from the two separate SPI buses as well as the IMU. Thread-safe first-in first-out queues with limited capacity (to avoid keeping old samples) are used to synchronize readings between threads and send them to the destination device.

D. Calibration

As shown in the modeling section and in previous work [30], the differential sensor output is very linear and the forces can be computed from linear equations. Thus, a linear calibration may, in theory, effectively convert from the raw Hall effect sensor readings to the applied forces and torques.

However, there are many potential non-linearities that are unmodeled. These include flex in the mechanical parts, non-linear bending and hysteresis of the suspension elements, saturation of the Hall effect sensors, shifting of the suspension or shell, and slight inconsistencies in the magnets or sensors and how they are mounted, among other factors. Thus, it is likely that a non-linear calibration would outperform the linear one, as has been shown in similar situations [39].

Therefore, both linear and non-linear calibrations were carried out and compared. The $2n$ raw Hall effect sensor measurements are input to the calibration rather than the difference and sum since this is more flexible and the calibration can learn to take the difference if needed. For the linear calibration, linear least squares was used to determine a 6×12 calibration matrix from raw voltages to force/torque. For non-linear calibration, neural networks were used. Hyperparameters of hidden layer sizes, number of hidden layers, activation functions, cost function, optimization algorithm, and momentum were varied to determine best architecture.

For the calibration tests, a jig was built to gather data. A 3D-printed dummy ultrasound probe was mounted directly to an ATI Nano25 sensor which was rigidly attached to a table. The sensing hardware was installed on the dummy probe as seen in Fig. 6. Both the master PCB and the ATI sensor were connected to a Windows PC, and samples were recorded with timestamps using a Python script. The data was aligned using the timestamps, and resampled to line up element-wise. For the numerical calibration, the ATI torques in Nm were first multiplied by 10 so their scale was the same as the forces. The input Hall effect sensor readings were scaled down by 2^{12} to fall between 0 and 1.

A random series of forces and torques was applied to the ultrasound probe while 20000 samples were recorded at 60 Hz. This was repeated three times, and the data sets combined. The 60000 samples were split 75%/25% into a training and a testing set, and 5-fold cross-validation was used on the training set for hyperparameter tuning. An additional measurement was taken without applying any forces, to study the noise characteristics, and a further set of forces and torques (20000 samples) was recorded after disassembly and reassembly to evaluate the calibration under potentially slightly different conditions.

IV. RESULTS

A. Suspension Tests

A variety of inner and outer diameters were tested for the suspension elements. In each case, the response was linear, with a much higher K value in axial compression/tension than in bending. The K value for each axis was determined by fitting a line and taking the slope. The results are outlined in Table I and Fig. 10.

Substituting these values into the simulation, the OD-ID combination of 6-4 mm gave the desired response. This was therefore fabricated in larger quantity and installed in the device.

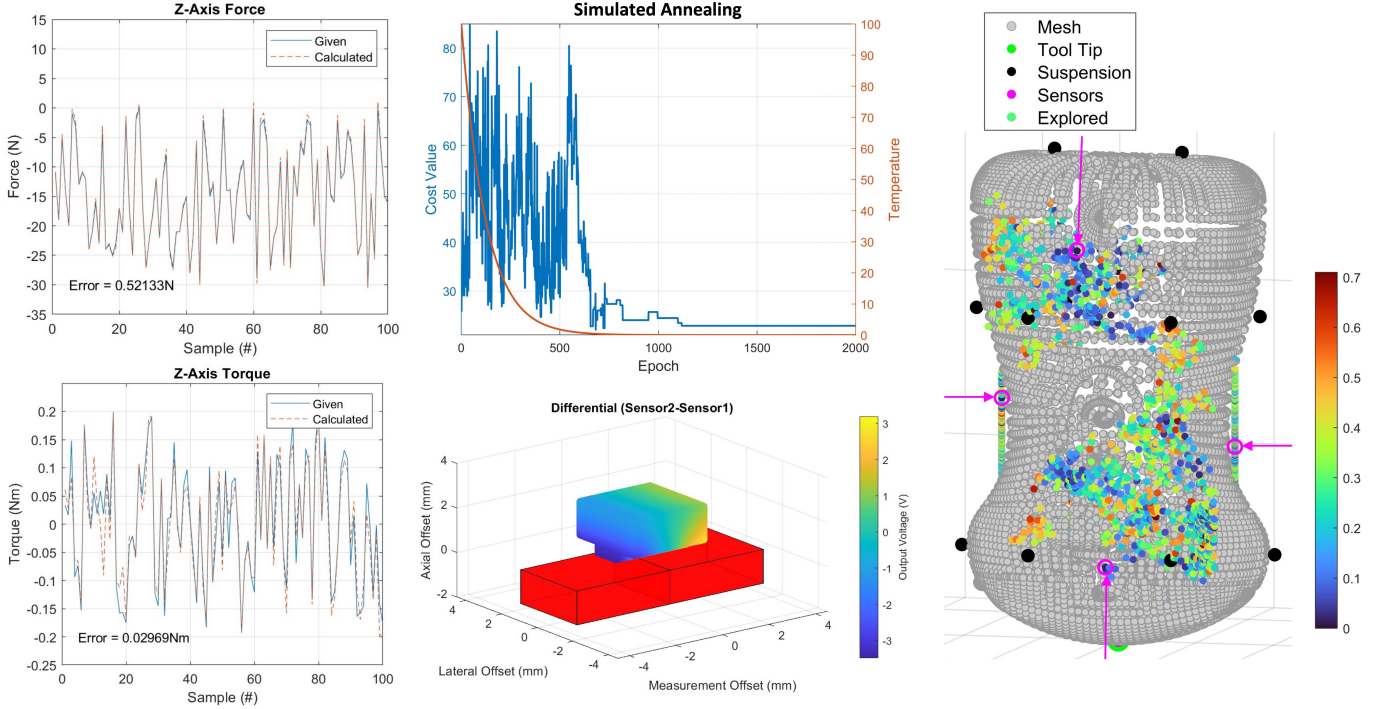


Fig. 9. Simulation results from simulated annealing for SM configuration optimization. The look-up-table for sensor output as a function of magnet displacement is shown on the bottom center, where each red box is a Hall effect sensor. The probe mesh with suspension locations (black) and explored sensor configurations (color map by cost value) are on the right. The final, optimal locations are shown in pink. The associated cost and annealing temperature are plotted in the top center, showing convergence to a low error value with an exponential cooling schedule. The simulation output and calibration with the optimal sensor configuration is shown on the left. The three axes of force and torque were similar, so only z is shown as an example. All errors are listed in Section IV-B.

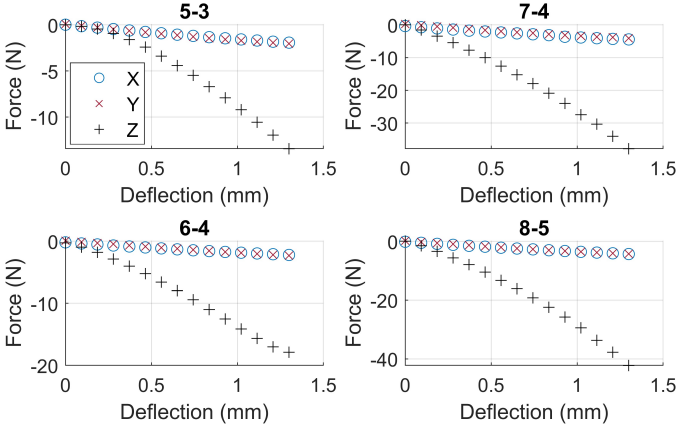


Fig. 10. Suspension test results from four of the five configurations in Table I. The responses are linear and the z axis (compression) is consistently much stiffer than x and y (bending). The plots are labeled OD-ID.

B. Simulation

The optimization procedure described in Section II-B was carried out with the mesh of a C3HD 3 ultrasound probe (Clarius, Vancouver, BC), and the chosen suspension elements. The suspension locations shown in Fig. 9 were used, and the SM locations limited to two per wide face and one per narrow face, for a total of six. The resulting optimal configuration is shown by the pink marks in Fig. 9, and the corresponding calibrated forces and torques are also shown. The RMS linear calibration error for force and torque were $[0.703, 1.39, 0.521]$ N and

TABLE I
SUSPENSION TEST RESULTS IN UNITS OF N/M. THE TORSIONAL STIFFNESS (K_θ) IS IN NMM/RAD.

OD-ID (mm)	K_{xx}	K_{yy}	K_{zz}	K_θ
5-3	1421	1571	11308	14.5
6-4	1472	1749	13335	24.9
8-5	2429	2580	27435	51.5
7-4	3035	3192	26321	34.1
8-3	4847	3834	42655	-

$[0.068, 0.083, 0.029]$ Nm respectively. In the Discussion section, these are compared to the values obtained in the real system, which are shown below.

C. Calibration

With a linear least squares calibration (performed on the training data), the RMS calibration error between our sensor and the ATI was (in N and Nm, on the testing data):

$$e_{RMS} = [1.04 \ 1.57 \ 2.06 \ 0.0918 \ 0.0753 \ 0.0545]$$

This is shown in Fig. 11, with the calibration applied to the additional measurement. Here x is to the right in Fig. 6, y is into the page, and z is up.

The tracking appears good, but the RMS error is relatively large. Instead, the neural network calibration was tested. In total, 44 different sets of hyperparameters were checked. It was quickly found that tanh activation significantly outperformed ReLU and logistic functions. An ADAM solver with adaptive

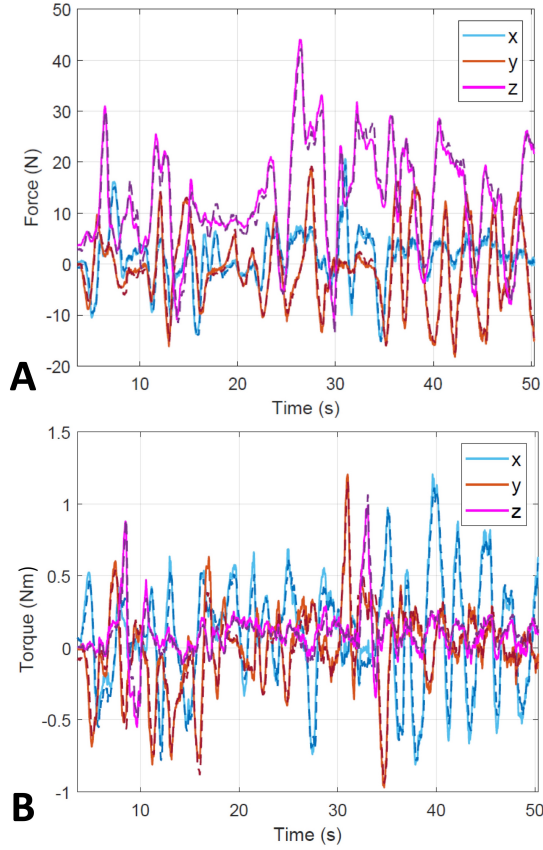


Fig. 11. Calibrated output from the DMFS sensor (dashed lines) compared to the ATI sensor (solid lines) using a linear calibration. (A) shows forces and (B) shows torques.

learning rate starting at 0.0005 with a Nesterov momentum of 0.9 worked best. L2 regularization was also used. With these settings, different numbers of hidden layers and layer sizes were tested. The results are shown in Fig. 12A. Using the elbow method, a 5-layer, 150 neuron-per-layer architecture was chosen. It outperformed the 4-layer 256-neuron model despite having fewer neurons. Models with 1, 2, or 3 hidden layers were clearly underfitting and hardly outperforming the linear calibration. The chosen architecture achieves RMS errors of (in N and Nm):

$$e_{train} = [0.463 \quad 0.506 \quad 0.532 \quad 0.0373 \quad 0.0360 \quad 0.0253]$$

$$e_{test} = [0.544 \quad 0.593 \quad 0.660 \quad 0.0432 \quad 0.0403 \quad 0.0271]$$

The similarity in performance between the training and testing data sets imply that the model is not overfitting much. This is discussed in Section V. The error is well under 1 N and 0.05 Nm in each axis compared to the ATI sensor. The RMS values of the dataset labels are:

$$y_{RMS} = [5.99 \quad 7.32 \quad 12.2 \quad 0.411 \quad 0.371 \quad 0.138]$$

which implies an approximate percent error of

$$\%e = [9.08 \quad 8.11 \quad 4.35 \quad 10.51 \quad 9.71 \quad 18.3]$$

Conversely, the range of the measured forces and torques was

$$\text{rng}(y) = \pm [48.9 \quad 42.6 \quad 65.8 \quad 2.64 \quad 2.49 \quad 1.66]$$

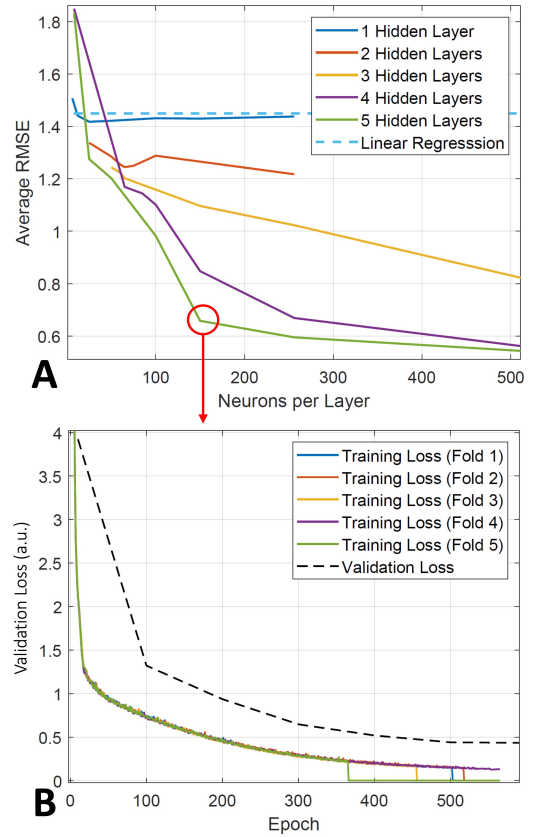


Fig. 12. (A) Neural network calibration RMS error, averaged over the five validation folds, versus hidden layer size with different numbers of hidden layers. (B) The learning curve of the chosen architecture, for both training and validation loss from the five folds.

so as a percentage of the range, the error was only

$$\%e_{rng} = [1.11 \quad 1.39 \quad 1.00 \quad 1.64 \quad 1.62 \quad 1.63]$$

As a further test, the calibration network was applied to the additional, separate measurement. This is plotted in Fig. 13.

D. Rate and Noise characteristics

The rate of SPI data transfer for one SM to the master board was determined to be 12 kHz with our firmware. Given that three SMs are read in series, the sampling rate is 4 kHz. Including synchronization, processing, and communication delays, the practical readout rate at which the sensor delivered samples over UART to the PC was reduced to 1276.8 ± 3.3 Hz on average. This is fast enough for haptics applications where a control loop of approximately 1 kHz is commonly employed.

From one measurement in which no forces were applied, we can analyze the noise characteristics of the sensor. The 12 raw sensor measurements are each normally distributed with an average standard deviation of 1.92 LSBs, or 0.047% of the full-scale range of 2^{12} . When the raw values are substituted into the calibration in the previous section, the standard deviation becomes 0.0043 N and 0.0314 Nmm. With a 95% confidence interval of 2σ , the force and torque resolution are 0.0086 N and 0.0628 Nmm respectively. Using a fast Fourier transform, the average noise power spectral density is




Fig. 13. Calibrated output from the DMFS sensor (dashed lines) compared to the ATI sensor (solid lines) using the neural network calibration applied to unseen data from the separate measurement. (Top) Shows all the forces. (Left) Shows a subset for better visualization, and (Right) shows a similar subset of the torques.

$8.867 \times 10^{-4} \text{ N}/\sqrt{\text{Hz}}$ and $1.267 \times 10^{-4} \text{ Nmm}/\sqrt{\text{Hz}}$ for force and torque respectively. The noise power spectral density is shown in Fig. 14.




Fig. 14. Noise power spectral density of the 6-DOF sensor

E. Sensor Module Analysis

In addition to the locations of the SMs, their orientation is also relevant to the effectiveness of the sensor. The results presented thus far were obtained using the optimal orientations found in Section IV-B. Referencing Fig. 6, the measurement axes of SMs 2 and 4 were vertical, SMs 1 and 3 were horizontal, and SM 5 and 6's measurement axes come out

of the page. Using standard ultrasound notation, these are axial, lateral, and elevational respectively. The latter two SMs should stay in this configuration because otherwise there are no sensors sensitive to the force direction out of the page. However, it is possible that rotating SMs 1, 2, 3, and/or 4 by 90° could improve the performance. Therefore, two further tests were carried out. First, SM1 and 2 were each rotated 90° so the sensors were asymmetric, with 1 and 4 vertical, and 2 and 3 horizontal. This is referred to as trial 1. Next, SM3 and 4 were each also rotated 90° so now 1 and 3 were vertical, and 2 and 4 were horizontal. This is trial 2. Trial 0 is the original configuration described above.

Performing the same neural network calibration procedure for each trial, the RMS validation errors in Table II were found. The original configuration of trial 0 clearly outperforms the other two configurations, showing that the optimization algorithm was effective in choosing the SM orientations, and also that the SM orientation makes a significant difference in the sensor performance

TABLE II
RMS VALIDATION ERROR (AVERAGED OVER THE x , y , AND z AXES)
VERSUS SM ORIENTATION. TRIAL 0, THE OUTPUT OF THE OPTIMIZATION
PROCEDURE, PERFORMS SIGNIFICANTLY BETTER THAN THE OTHER TWO.

Trial	Horiz.	Vert.	Force (N)	Torque (Nm)
0	SM1,3	SM2,4	0.582	0.042
1	SM2,3	SM1,4	0.801	0.050
2	SM2,4	SM1,3	0.792	0.053

F. Shell Rigidity

The rigidity of the shell was tested using FEA by placing elastic supports with the known suspension stiffness values in their chosen positions. A 25 N force, the maximum experienced in typical ultrasound procedures, was applied to the right side of the shell. The setup is shown in Fig. 15. Comparing the deflection on the left and right sides of the shell shows how much the shell itself deformed versus how much the suspension moved. On average, the shell deflection was 0.083 mm while the suspension deflection was 0.95 mm. Thus, shell deformation constitutes 7.8% of the total deflection.




Fig. 15. Finite element analysis of outer shell rigidity. (Left) Shell showing elastic supports (blue) and applied force (pink). (Right) Shell colored according to deflection, with several point measurements made symmetrically on both sides. The difference between the two sides shows the shell deflection due to the maximum design force of 25 N. It is on average 0.08 mm.

V. DISCUSSION

Comparing the results from simulation in Section IV-B and the real system in Section IV-C, we see very similar errors. Indeed, the real force and torque errors from the linear calibration are on average within 0.68 N and 0.019 Nm of the simulated values, which also used a linear calibration. Moreover, the optimal sensor locations are approximately where one would intuitively place them, and the chosen sensing axes corresponded with the ones that experimentally performed best. Thus, the simulation is realistic and the optimization is effective. It is possible that other orientations, not just vertical or horizontal, could improve performance. Investigation of this is left for future work.

Given the different SM orientations, the question arises which SMs are actually used for forces and torques in which axis. If only forces in one direction were relevant, some sensors could be removed, or they could all be oriented in the same direction. For example, in the current configuration, only SMs 2 and 4 are sensitive to vertical (z) forces. For x and y forces, two SMs are directly sensitive to the force and two SMs have the magnet move axially closer or farther away from the Hall effect sensors. This could be why the error in z is slightly higher than in the other two axes. Conversely, for torques, each axis of load excites responses in four SMs. However,

these are all highly coupled between axes, with overlapping responses from the SMs for different torques. Thus, the system is far from being linearly separable, and the more complex architecture of the neural network is warranted.

However, there is always a risk of overfitting. According to Heaton, it is rarely necessary to use more than one hidden layer as this can be a universal approximator [47]. He further gives a rule of thumb that the number of hidden neurons should be between the size of the input and output layers. In this case, that would be about 8 neurons. Hadi et al. used a single hidden layer with just 5 neurons to good effect [39]. However, from Fig. 12, a single hidden layer architecture with 5 or 8 neurons barely outperformed the linear calibration. Conversely, Hagan states that one can usually avoid overfitting by using

$$n_{hid} \leq \frac{n_{samples}}{\alpha(N_{in} + N_{out})}$$

hidden neurons [48]. In this case, $n_{samples} = 45000$, $N_{in} = 12$, and $N_{out} = 6$. The factor α should be between 2 and 5. Using these values, less than approximately 500 to 1250 hidden neurons should not lead to overfitting. Indeed, the chosen architecture uses 750 neurons, which is unlikely to overfit. The next-best choice was a 4-layer network with 256 neurons per layer, which exceeds 1000 neurons and actually achieved worse results. Slight improvement in test error could be gained using 5 layers and 256 neurons per layer, but this is likely overfitting and the performance improvement is marginal. Having different numbers of neurons in each layer was also tested, for example having [50, 150, 256, 150, 50], but the performance was consistently worse. With the chosen architecture (5 hidden layers, 150 neurons per layer), the validation and test errors are similar to the training error, indicating that the network generalizes well. In comparison, architectures with fewer hidden layers were clearly underfitting. Regularization was also used, which limits overfitting. Hence, the chosen architecture is justified and is most likely not overfitting.

An additional point of concern for instrumenting manual tools and devices is the added weight. The inner scaffold with magnets and suspension elements weighs 31.17 g while the full shell with all electronics shown in Fig. 6 weighs 202.98 g. In total, this amounts to 234.15 g. New PCBs have been designed but not yet tested which are substantially smaller (45% reduction in surface area). These also use a single flexible PCB with shared data, clock, power, and ground buses instead of the separate wires in Fig. 6. This will decrease the weight significantly. A rendering of the updated circuitry and integration is shown in Fig. 16.

Further improvements to the system are also possible. The ESP32 processor could be replaced by a field programmable gate array (FPGA) chip configured as six cores, so each SM could be processed in parallel for minimum latency. Such a system is described by Hadi et al. for optical force sensing [31]. Other hardware modifications could be made, especially in reducing manufacturing tolerances and exploring stiffer materials, to obtain more accurate measurements. Currently, approximately 7.8% of the total deflection is used in deforming the shell rather than the suspension. Some of this nonlinearity




Fig. 16. Outer shell with updated PCBs. The master board is miniaturized and has a button on the opposite face for resetting. The messy wiring is replaced by a flexible PCB.

is likely captured by the neural network calibration, but a stiffer material may improve accuracy. AMR sensors could also be explored, despite the added circuitry which would increase the size of the sensor modules. In terms of firmware, for wireless integration into systems, Bluetooth Classic and/or UDP communication over Wi-Fi should be explored. In the latter case, the sensor must connect to a Wi-Fi network and know the IP address of the destination device. To avoid having to update the firmware manually for every change of network or device, an initial automatic connection and exchange of connectivity information over BLE could be used. This is in fact already implemented, but the UDP communication part will be tested in future work.

This paper has introduced a sensor modality and optimization-based design procedure which was tested on one specific ultrasound probe. Integration into other point-of-care ultrasound devices using the same approaches will constitute future work. Beyond ultrasound, many other applications can also be explored. The Hall effect sensors and signal processing circuitry could be packaged into a single, very small integrated circuit which could enable much lower-profile hardware integration and even a small standalone force/torque sensor.

The benefits of this sensing modality are that it is modular, non-contact, and thus simple to integrate into a variety of tools without making them bulky or difficult to assemble. This enables practical use in medical environments, where the tool can be removed and cleaned or replaced. The sensor itself can be disinfected if the PCBs are sealed with a very thin layer of casting resin. The size and weight are low, and the flexibility enables the use of optimal design principles. Furthermore, the mechanical components can all be 3D-printed. This cost is negligible since integrating off-the-shelf sensors also involves fabricating a mechanical shell or structure. The electrical components are inexpensive and require no special preparation or treatment, so the whole integration of our sensor can be completed at a small fraction of the cost of existing multi-

axis sensors. This may be important for remote guidance in low-resource environments. Despite this, the sensor still achieves good accuracy and low noise. Moreover, the PCB design facilitates flexible wired or wireless integration into a variety of systems.

VI. CONCLUSION

This paper has introduced a novel force sensing modality and optimization-based design and integration procedure for the instrumentation of manual tools. Made from inexpensive electrical components and 3D-printed with no special fabrication procedures or delicate mechanisms, this sensor is low-cost and easy to reproduce for a variety of applications including remote guidance of people and autonomous or teleoperated robotics. Integration into complex systems is simple through USB, Bluetooth, or Wi-Fi communication with a practical readout rate of 1.3 kHz over USB. Despite the simplicity, the sensor achieves good performance through a neural network-based calibration, with average root-mean-square test errors of 1.17% and 1.63% of the full-scale range for force and torque respectively. The sensor has low noise and a 95% confidence interval resolution of 0.0086 N and 0.0628 Nm. This will hopefully enable more effective tele-guidance, robotics, and autonomy through force feedback for manual tasks.

ACKNOWLEDGEMENTS

We gratefully acknowledge scholarship support from the Vanier Canada Graduate Scholarships program, infrastructure support from CFI and funding support from NSERC and the Charles Laszlo Chair in Biomedical Engineering, as well as infrastructure, technical, and funding support from Rogers Communications and MITACS.

APPENDIX I

A coordinate-free explanation of the system model is described in Section II-B. However, to implement this in practice, coordinate transforms must be taken into account. We rewrite Equation 1 in the coordinate frame of the individual suspension element:

$$\vec{w}_j = K_j \vec{\delta}_j \quad (10)$$

where K_j was measured in Section IV-A. The suspension members are positioned around the tool to ensure relatively uniform and symmetric suspension. Each member has a position \vec{x}_j and rotation $R_j \in \text{SO}(3)$ relative to the coordinate frame of the tool. Thus, the member pose is a homogeneous transform from the tool frame:

$$T_i = \begin{bmatrix} R_i & \vec{x}_i \\ 0 & 1 \end{bmatrix}$$

Equation 10 is written in the member's local coordinate system. However, the forces applied to the tool are known in the tool coordinates. Hence, we must convert Equation 10 to the tool frame. Define

$$\tilde{R}_j = \begin{bmatrix} R_j & 0 \\ 0 & R_j \end{bmatrix}$$

The stiffness matrix of each suspension element in the tool frame is then

$$K_j = \tilde{R}_j K \tilde{R}_j^\top \quad (11)$$

where K is the same for every member.

Let us additionally examine Equation 2. Suspension member j experiences translational deflection due to the rotation at the tool-tip and the offset of the member from the tool-tip. Thus,

$$\vec{\delta}_j = \begin{bmatrix} \vec{\delta}_t + (R(\vec{\delta}_\theta) \vec{x}_j - \vec{x}_j) \\ \vec{\delta}_\theta \end{bmatrix} \quad (12)$$

Where $R(\vec{\delta}_\theta) = e^{[\vec{\delta}_\theta]_\times}$ is the rotation matrix representation of the rotational deflection. However, the rotations $\vec{\delta}_\theta$ are very small, so we can take the first two terms of the Taylor expansion for $R(\vec{\delta}_\theta) = e^{[\vec{\delta}_\theta]_\times}$:

$$R_t(\vec{\delta}_t) \approx I_3 + [\vec{\delta}_\theta]_\times$$

Substituting this into Equation 12, we arrive at Equation 2. Now we can write Equation 4 concretely in the tool's coordinate system, to be solved numerically:

$$W = \sum_{j=1}^m \begin{bmatrix} I_3 & 0_3 \\ [\vec{x}_j]_\times & I_3 \end{bmatrix} \tilde{R}_j K \tilde{R}_j^\top \begin{bmatrix} \vec{\delta}_t + [\vec{\delta}_\theta]_\times \vec{x}_j \\ \vec{\delta}_\theta \end{bmatrix} \quad (13)$$

Note, using Equation 2 causes the simulation to run approximately four times faster than when using Equation 12.

VII. APPENDIX II

A. Simulated Annealing

One possible search algorithm for the optimization is simulated annealing. This uses a “temperature” parameter, T , which describes the randomness of the search and gradually “cools”, or decreases in randomness, according to some temperature schedule as the search progresses, in analogy to annealing used in metallurgy. The optimization is described in Algorithm 2.

Algorithm 2 Simulated Annealing Optimization

Randomly initialize the set of sensor positions and orientations, $S \in \mathcal{D}$

Choose the desired number of epochs, Q

Choose the desired search radius, σ

Define $w(S)$ from Equation 8

$T \leftarrow \text{logspace}(2, -4, Q)$

for $t \in T$ **do**

 Perform the simulation to find $w(S)$

 Choose nearby sensor positions, $S' \sim \mathcal{N}(S, \sigma)$, $S' \in \mathcal{D}$

 Perform the simulation to find $w(S')$

if $w(S') < w(S)$ **then**

$S \leftarrow S'$

else

 Sample p from $\text{Unif}[0, 1]$

if $\exp(-(w(S') - w(S))/t) \geq p$ **then**

$S \leftarrow S'$

end if

end if

end for

The temperature schedule used in Algorithm 2 is an exponential cooling from 10^2 to 10^{-4} in Q steps, using the MATLAB logspace syntax. Different cooling schedules can be explored, and would tend to give slightly different results.

This method is an improvement on a brute force search of the entire space because as the temperature cools, it converges on the optimum and becomes a more greedy search. However, the stochastic nature greatly outperforms a simple greedy search by ensuring the solution space is well explored and the algorithm probably does not stop in a local optimum.

REFERENCES

- [1] S. E. Salcudean, H. Moradi, D. Black, and N. Navab, “Robot-assisted medical imaging: A review,” *Proceedings of the IEEE*, 2022.
- [2] S. E. Salcudean, G. Bell, S. Bachmann, W.-H. Zhu, P. Abolmaesumi, and P. D. Lawrence, “Robot-assisted diagnostic ultrasound—design and feasibility experiments,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 1062–1071, Springer, 1999.
- [3] M. Akbari, J. Carriere, T. Meyer, R. Sloboda, S. Husain, N. Usmani, and M. Tavakoli, “Robotic ultrasound scanning with real-time image-based force adjustment: Quick response for enabling physical distancing during the covid-19 pandemic,” *Frontiers in Robotics and AI*, vol. 8, 2021.
- [4] C. Delgorge, F. Courrèges, L. A. Bassit, C. Novales, C. Rosenberger, N. Smith-Guerin, C. Brù, R. Gilabert, M. Vannoni, G. Poisson, *et al.*, “A tele-operated mobile ultrasound scanner using a light-weight robot,” *IEEE transactions on information technology in biomedicine*, vol. 9, no. 1, pp. 50–58, 2005.
- [5] K. Mathiassen, J. E. Fjellin, K. Glette, P. K. Hol, and O. J. Elle, “An ultrasound robotic system using the commercial robot ur5,” *Frontiers in Robotics and AI*, vol. 3, p. 1, 2016.
- [6] D. Black, Y. Oloumi Yazdi, A. H. Hadi Hosseinabadi, and S. Salcudean, “Human teleoperation—a haptically enabled mixed reality system for teleultrasound,” *Human-Computer Interaction*, pp. 1–24, 2023.
- [7] D. Black and S. Salcudean, “Mixed reality human teleoperation,” in *2023 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 375–383, IEEE, 2023.
- [8] O. Pallotta and A. Roberts, “Musculoskeletal pain and injury in sonographers, causes and solutions,” *Sonography*, vol. 4, no. 1, pp. 5–12, 2017.
- [9] A. E. Drake, J. Hy, G. A. MacDougall, B. Holmes, L. Icken, J. W. Schrock, and R. A. Jones, “Innovations with tele-ultrasound in education sonography: the use of tele-ultrasound to train novice scanners,” *The ultrasound journal*, vol. 13, no. 1, pp. 1–8, 2021.
- [10] N. J. Soni, J. S. Boyd, G. Mints, K. C. Proud, T. P. Jensen, G. Liu, B. K. Mathews, C. K. Schott, L. Kurian, C. M. LoPresti, *et al.*, “Comparison of in-person versus tele-ultrasound point-of-care ultrasound training during the covid-19 pandemic,” *The ultrasound journal*, vol. 13, no. 1, pp. 1–7, 2021.
- [11] H. Sharma, L. Drukker, A. T. Papageorghiou, and J. A. Noble, “Multi-modal learning from video, eye tracking, and pupillometry for operator skill characterization in clinical fetal ultrasound,” in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 1646–1649, IEEE, 2021.
- [12] T. Kaneko, N. Kagiya, Y. Nakamura, T. Hirasawa, A. Murata, R. Morimoto, S. Miyazaki, and T. Minamino, “Effectiveness of real-time tele-ultrasound for echocardiography in resource-limited medical teams,” *Journal of Echocardiography*, vol. 20, no. 1, pp. 16–23, 2022.
- [13] K. Jemal, D. Ayana, F. Tadesse, M. Adefris, M. Awol, M. Tesema, B. Dagne, S. Abeje, A. Bantie, M. Butler, *et al.*, “Implementation and evaluation of a pilot antenatal ultrasound imaging programme using tele-ultrasound in ethiopia,” *Journal of Telemedicine and Telecare*, p. 1357633X221115746, 2022.
- [14] Z. Jiang, M. Grimm, M. Zhou, Y. Hu, J. Esteban, and N. Navab, “Automatic force-based probe positioning for precise robotic ultrasound acquisition,” *IEEE Transactions on Industrial Electronics*, vol. 68, no. 11, pp. 11200–11211, 2020.
- [15] S. Virga, O. Zettinig, M. Esposito, K. Pfister, B. Frisch, T. Neff, N. Navab, and C. Hennersperger, “Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms,” in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 508–513, IEEE, 2016.

- [16] L. Drukker, H. Sharma, R. Droste, M. Alsharid, P. Chatelain, J. A. Noble, and A. T. Papageorghiou, "Transforming obstetric ultrasound into data science using eye tracking, voice recording, transducer motion and ultrasound video," *Scientific Reports*, vol. 11, no. 1, p. 14109, 2021.
- [17] R. Droste, L. Drukker, A. T. Papageorghiou, and J. A. Noble, "Automatic probe movement guidance for freehand obstetric ultrasound," in *Medical Image Computing and Computer Assisted Intervention-MICCAI 2020: 23rd International Conference, Lima, Peru, October 4-8, 2020, Proceedings, Part III 23*, pp. 583-592, Springer, 2020.
- [18] T. Schimmoeller, R. Colbrunn, T. Nagle, M. Lobosky, E. E. Neumann, T. M. Owings, B. Landis, J. E. Jelovsek, and A. Erdemir, "Instrumentation of off-the-shelf ultrasound system for measurement of probe forces during freehand imaging," *Journal of biomechanics*, vol. 83, pp. 117-124, 2019.
- [19] M. W. Gilbertson and B. W. Anthony, "An ergonomic, instrumented ultrasound probe for 6-axis force/torque measurement," in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 140-143, IEEE, 2013.
- [20] M. O. Harris-Love, C. Ismail, R. Monfaredi, H. J. Hernandez, D. Pennington, P. Woletz, V. McIntosh, B. Adams, and M. R. Blackman, "Interrater reliability of quantitative ultrasound using force feedback among examiners with varied levels of experience," *PeerJ*, vol. 4, p. e2146, 2016.
- [21] G. P. Mylonas, P. Giataganas, M. Chaudery, V. Vitiello, A. Darzi, and G.-Z. Yang, "Autonomous efast ultrasound scanning by a robotic manipulator using learning from demonstrations," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3251-3256, IEEE, 2013.
- [22] Q. Huang, J. Lan, and X. Li, "Robotic arm based automatic ultrasound scanning for three-dimensional imaging," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 2, pp. 1173-1182, 2019.
- [23] Y. Wei and Q. Xu, "An overview of micro-force sensing techniques," *Sensors and Actuators A: Physical*, vol. 234, pp. 359-374, 2015.
- [24] J. O. Templeman, B. B. Sheil, and T. Sun, "Multi-axis force sensors: A state-of-the-art review," *Sensors and Actuators A: Physical*, vol. 304, p. 111772, 2020.
- [25] S. A. Liu and H. L. Tzo, "A novel six-component force sensor of good measurement isotropy and sensitivities," *Sensors and Actuators A: Physical*, vol. 100, no. 2-3, pp. 223-230, 2002.
- [26] L. Xiong, G. Jiang, Y. Guo, and H. Liu, "A three-dimensional fiber bragg grating force sensor for robot," *IEEE Sensors Journal*, vol. 18, no. 9, pp. 3632-3639, 2018.
- [27] Y.-L. Park, B.-R. Chen, and R. J. Wood, "Design and fabrication of soft artificial skin using embedded microchannels and liquid conductors," *IEEE Sensors journal*, vol. 12, no. 8, pp. 2711-2718, 2012.
- [28] C. Lebossé, P. Renaud, B. Bayle, and M. de Mathelin, "Modeling and evaluation of low-cost force sensors," *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 815-822, 2011.
- [29] A. H. H. Hosseinabadi, M. Honarvar, and S. E. Salcudean, "Optical force sensing in minimally invasive robotic surgery," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 4033-4039, IEEE, 2019.
- [30] D. Black, A. H. H. Hosseinabadi, N. R. Pradnyawira, M. Pol, M. Nogami, and T. Salcudean, "Towards differential magnetic force sensing for ultrasound teleoperation," in *2023 IEEE World Haptics Conference (WHC)*, pp. 333-339, IEEE, 2023.
- [31] A. H. H. Hosseinabadi, D. Black, and S. E. Salcudean, "Ultra low-noise fpga-based six-axis optical force-torque sensor: Hardware and software," *IEEE Transactions on Industrial Electronics*, vol. 68, no. 10, pp. 10207-10217, 2020.
- [32] B.-Y. Lee, J. Kim, H. Kim, C. Kim, and S.-D. Lee, "Low-cost flexible pressure sensor based on dielectric elastomer film with micro-pores," *Sensors and Actuators A: Physical*, vol. 240, pp. 103-109, 2016.
- [33] D. Jones, L. Wang, A. Ghanbari, V. Vardakastani, A. E. Kedgley, M. D. Gardiner, T. L. Vincent, P. R. Culmer, and A. Alazmani, "Design and evaluation of magnetic hall effect tactile sensors for use in sensorized splints," *Sensors*, vol. 20, no. 4, p. 1123, 2020.
- [34] P. J. Kyberd and P. H. Chappell, "A force sensor for automatic manipulation based on the hall effect," *Measurement Science and Technology*, vol. 4, no. 3, p. 281, 1993.
- [35] P. Yu, X. Qi, W. Liu, and X. Fu, "Development of a compliant magnetic 3-d tactile sensor with amr elements," in *Intelligent Robotics and Applications: 7th International Conference, ICIRA 2014, Guangzhou, China, December 17-20, 2014, Proceedings, Part II 7*, pp. 484-491, Springer, 2014.
- [36] A. Arami, M. Simoncini, O. Atasoy, S. Ali, W. Hasenkamp, A. Bertsch, E. Meurville, S. Tanner, P. Renaud, C. Dehollain, *et al.*, "Instrumented knee prosthesis for force and kinematics measurements," *IEEE transactions on automation science and engineering*, vol. 10, no. 3, pp. 615-624, 2013.
- [37] H. Kristanto, P. Sathe, A. Schmitz, T. P. Tomo, S. Somlor, and S. Sugano, "A wearable three-axis tactile sensor for human fingertips," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4313-4320, 2018.
- [38] S. D. M. Nasab, A. Beiranvand, M. T. Masouleh, F. Bahrani, and A. Kalhor, "Design and development of a multi-axis force sensor based on the hall effect with decouple structure," *Mechatronics*, vol. 84, p. 102766, 2022.
- [39] A. Hadi-Hosseinabadi and S. Salcudean, "Multi-axis force sensing in robotic minimally invasive surgery with no instrument modification," *arXiv preprint arXiv:2103.11116*, 2021.
- [40] A. Bicchi, "A criterion for optimal design of multi-axis force sensors," *Robotics and Autonomous Systems*, vol. 10, no. 4, pp. 269-286, 1992.
- [41] Y. Hou, D. Zeng, J. Yao, K. Kang, L. Lu, and Y. Zhao, "Optimal design of a hyperstatic stewart platform-based force/torque sensor with genetic algorithms," *Mechatronics*, vol. 19, no. 2, pp. 199-204, 2009.
- [42] E. W. Cheney and D. R. Kincaid, *Numerical mathematics and computing*. Cengage Learning, 2012.
- [43] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671-680, 1983.
- [44] D. Black and S. Salcudean, "Robust object pose tracking for augmented reality guidance and teleoperation," *TechRxiv*, 2023.
- [45] D. Black and S. Salcudean, "A mixed reality system for human teleoperation in tele-ultrasound," pp. 91-92, Hamlyn Symposium for Medical Robotics, July 2022.
- [46] D. Black and S. Salcudean, "Human-as-a-robot performance in mixed reality teleultrasound," *International Journal of Computer Assisted Radiology and Surgery*, pp. 1-8, 2023.
- [47] J. Heaton, *Introduction to neural networks with Java*. Heaton Research, Inc., 2008.
- [48] M. T. Hagan, H. B. Demuth, and M. Beale, *Neural network design*. PWS Publishing Co., 1997.