Head-On: Robust Tactile Force Sensing for Quadrupedal Interaction

Stergios E. Bachoumas, John Cahill and Panagiotis Artemiadis* *IEEE Senior Member*

***Abstract*— Quadrupedal robots have achieved remarkable mobility and robustness in unstructured terrains, yet their capacity for deliberate physical interaction remains notably limited. Current methodologies for force and tactile sensing on quadrupeds are almost exclusively focused on the feet for estimating ground reaction forces, leaving the bulk of the robot's body, specifically the head, as an underutilized resource for purposeful environmental manipulation. To bridge this gap, we introduce *Head-On*: a novel multi-sensor tactile helmet architecture that re-purposes the quadruped's head into an active contact interface for forceful interaction with the environment. To enable these interactions, the *Head-On* architecture utilizes soft, magnetometer-based tactile sensors arranged in an optimized spatial configuration. To mitigate the severe kinematic and geomagnetic signal drift inherently induced by deploying magnetic sensors on dynamic, floating-base platforms, we propose a locomotion-robust, two-step calibration framework. This framework utilizes a causal one-dimensional Convolutional Neural Network (1D-CNN) for real-time contact estimation coupled with a dynamic re-baselining strategy, which continuously isolates no-contact phases to eliminate accumulated drift and anchor the sensor data to a valid operational domain. Experimental validation via controlled pushing tasks demonstrates that our approach successfully eliminates the unmitigated drift that otherwise fundamentally corrupts force estimation. Coupled with a lightweight Long Short-Term Memory (LSTM) and Multi-Layer Perceptron (MLP) neural network architecture, the system achieves high-fidelity 3D force mapping with a Mean Absolute Error (MAE) of 8 N, closely tracking ground-truth measurements across extended, repeated physical interactions. By transforming the quadruped's head into a robust interface for physical interaction, this work empowers legged robots to evolve beyond pure locomotion, paving the way for advanced whole-body manipulation and collaborative tasks in complex, real-world environments.**

I. Introduction

Quadrupedal robots have seen remarkable advancements in recent years [1], demonstrating unprecedented agility and robustness in complex, unstructured terrains. However, the vast majority of research in this domain has been predominantly focused on locomotion, optimizing gait patterns, maintaining dynamic balance, and navigating uneven terrain [2], [3]. While these achievements have established quadrupeds as highly capable platforms for mobility and inspection, their capacity to intentionally and physically interact with their surroundings remains notably limited.

*This material is based upon work supported by the National Science Foundation under Grant No. #2014264, as well as by the U.S. Air Force Office of Scientific Research (AFOSR) award FA9550-18-1-0221.

Stergios E. Bachoumas, John Cahill and Panagiotis Artemiadis are with the Mechanical Engineering Department, at the University of Delaware, Newark, DE 19716, USA. stevbach@udel.edu, jccahill@udel.edu, partem@udel.edu

*Corresponding author

When deliberate physical interaction is addressed in the literature, current methodologies rely heavily on the robot's extremities. Specifically, previous work that integrates force and tactile sensors on quadrupeds is almost exclusively focused on the feet of the robot [4]. These foot-mounted sensors are vital for estimating ground reaction forces and improving state estimation during walking, but they are not designed to facilitate deliberate, forceful interaction with objects, obstacles, or humans in the environment. Consequently, the bulk of the robot's body and specifically the head remains an underutilized resource for purposeful, physical environmental manipulation.

To bridge this gap, we introduce a novel paradigm for quadrupedal interaction that utilizes the robot's head as a primary, active contact interface. In this paper, we present *Head-On*: a tactile helmet for forceful quadruped interaction, seen in Fig. 1. This novel sensorized helmet utilizes tactile sensors to measure the physical interactions between the robot's head and its surroundings, allowing the quadruped to safely and deliberately push, nudge, and engage with the environment and human collaborators, transforming it from a purely mobile observer into an active physical agent. By integrating this technology into a novel morphology, *Head-On* capitalizes on the recent surge of interest in tactile sensing [5]–[7], proving its viability beyond traditional robotic arms and into whole-body quadrupedal interaction.

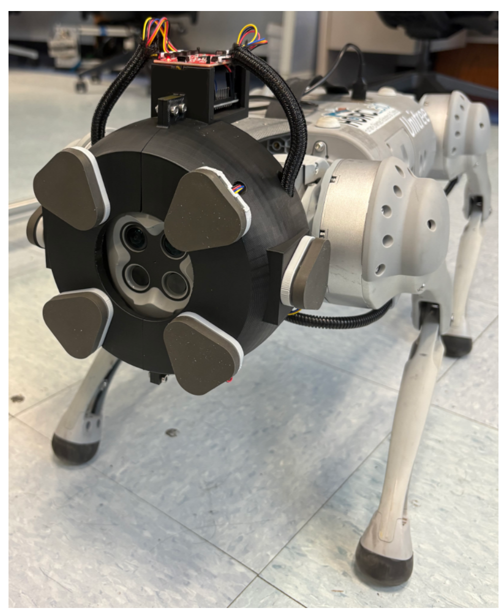

Fig. 1: Head-on: A tactile helmet that enables quadruped robots to interact with their environment, shown attached to the Unitree Go1 quadruped robot.

Realizing this capability required overcoming several hardware integration challenges to ensure the head could serve as a reliable and safe contact surface. The underlying tactile sensing system was subject to four strict design criteria. First, the sensing system needed to be low-cost to enable replication and scalability across multiple experimental configurations. Second, the system needed to be plug-and-play, allowing straightforward electrical integration and rapid deployment without extensive hardware modification to the base robot. Third, the sensors had to be compact and low-profile to avoid interfering with the robot's cameras, joints, or overall mobility. Finally, the sensing elements needed to be deployable as a unified contact surface, allowing multiple sensing elements to function together as a single large sensing interface. These constraints directly guided the evaluation of candidate tactile sensing technologies and the final design of the *Head-On* helmet.

This work advances the field of mobile physical interaction through three primary contributions. First, we introduce *Head-On*, a novel multi-sensor tactile helmet architecture that repurposes the quadruped's head as an active contact interface for whole-body environmental manipulation. Second, we present a locomotion-robust contact estimation algorithm that enables the use of magnetometer-based tactile sensors on free-base mobile platforms. By decoupling environmental contact forces from the magnetic and kinematic noise inherent to quadrupedal movement, this framework ensures reliable force estimation in unstructured settings. Finally, we provide preliminary real-world trials in select static and dynamic scenarios, suggesting the potential for controlled physical interaction while the robot is in motion.

The remainder of this paper is structured as follows: in Section II, we provide a full description of the methods used to design the helmet. In Section III, the optimization strategy for the shape of the helmet is introduced. In Section IV, the calibration procedure for the tactile sensors is discussed. In Section V, we present the results from our experiments, demonstrating the effectiveness of the helmet. Finally, in Section VI, we conclude with a summary of our findings and suggest directions for future research.

II. Materials and Methods

A. Robot Platform

The tactile sensing system, shown in Fig. 1, was developed for integration with the Unitree Go1 quadruped, allowing the robot to support controlled physical interaction with its environment. Because the robot is a mobile platform, several practical constraints governed the choice of the tactile sensor. We note here that the robot was chosen solely based on its widespread adoption as a low-cost research platform. Given the appropriate setup, the method presented herein could be applied to any quadruped robot, among other mobile platforms.

B. AnySkin Tactile Sensor

For the tactile sensors on the helmet, the AnySkin tactile sensor [8] was selected as the sensing modality for the unified

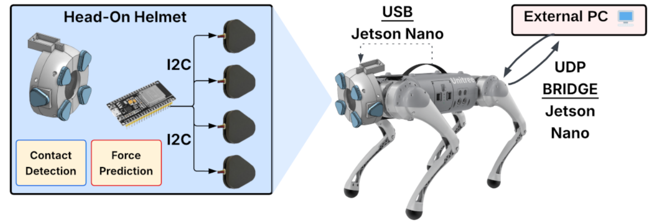

Fig. 2: Overview of the proposed *Head-on* helmet. The custom helmet is fitted on the Unitree Go 1 robot, and its electronics are powered by and communicate with one of the onboard computers of the robot via USB. A user can use a custom UDP bridge to stream the data from the helmet in the local network switch of the robot.

contact interface. AnySkin builds upon the same magnetic deformation principle used in ReSkin [9] but introduces fabrication and design improvements aimed at increasing measurement stability and usability. To keep the cost within a reasonable limit but without sacrificing sensing coverage, we decided to mount 4 AnySkin sensors on the face of the helmet. Each AnySkin unit consists of a permanently magnetized elastomer layer positioned above an array of embedded magnetometers. When the elastomer deforms under external loading, the internal magnetic field distribution changes. These variations are measured by the magnetometers.

C. Helmet Design

To build an effective helmet to complement the tactile sensing system, the helmet was designed as two parts that satisfy distinct operational requirements.

1) Helmet Base: The first part of the helmet, its base, was designed and manufactured with the mechanical structure of the helmet in mind. The helmet needed to conform seamlessly to the existing geometry of the Unitree Go1 head, ensuring zero occlusion or interference with the robot's onboard depth cameras, vision sensors, and kinematic workspace. Second, the mounting mechanism required a design that was both easily removable for rapid iteration and mechanically robust enough to remain stable under the shear and compressive forces of deliberate interaction with the environment. Finally, the proximity of the helmet to the quadruped's actuators necessitated careful consideration of signal integrity. Since the AnySkin tactile sensors utilize magnetometers, the helmet's layout was explicitly tailored to isolate the sensing elements and mitigate magnetic interference generated by the Go1's motors, ensuring high-fidelity data acquisition even under heavy load.

2) Sensing Face: The second part of the helmet, its face, was designed to maximize the data acquisition capabilities of the tactile sensors during robot interactions with its environment. To this end, we carefully designed an optimization pipeline for the shape of the face and evaluated 8 potential designs. The results of this optimization process are detailed in Section III.

D. Electrical Components & Communication Architecture

To support simultaneous multi-sensor operation while minimizing wiring complexity, we developed a centralized architecture using an ESP32 microcontroller (MCU) [10] as the primary data aggregation node. As illustrated in Fig. 2, the MCU communicates with all tactile units via a shared I²C bus, leveraging the native interface of the AnySkin electronics to enable direct integration without PCB modifications. This configuration streamlines the hardware footprint by multiplexing all sensor streams into a single high-level output interface. To resolve address conflicts among the identical AnySkin magnetometers, a TCA9548A 8-channel multiplexer was integrated to isolate each sensor on independent I²C channels. This hardware solution allows the ESP32 to cycle through sensors sequentially, maintaining a shared bus architecture without communication collisions.

The custom ESP32 firmware automatically detects and initializes active sensors at startup. During operation, the microcontroller sequentially queries each multiplexer channel, aggregates the tactile measurements, and streams the combined data via a single serial connection to the robot's onboard computer. As illustrated in Fig. 2, this onboard computer supplies power to the helmet electronics and handles high-level data routing, establishing a streamlined, centralized acquisition pipeline.

III. Helmet Shape Optimization

A. Optimization Setup

While the AnySkin sensors offer modular "plug-and-play" integration, their efficacy is fundamentally constrained by the geometry of the helmet's contact interface. Because tactile sensors require physical engagement to generate a signal, the sensing face must be engineered to maximize the probability and quality of environmental contact. This design challenge is exacerbated by the quadruped's high-degree-of-freedom locomotion; the shifting orientations and inertial forces of a dynamic platform render random placement insufficient. Consequently, the helmet's morphology and the sensors' placement must be optimized to ensure that the sensors remain responsive across the robot's full range of motion, transforming the surface from a passive mount into a strategic interface for data acquisition.

First, since the number of sensors mounted on the helmet is even, the sensors need to be placed in a symmetric pattern to reduce excessive turning torques about the point of contact. With this symmetry in mind, two patterns were of interest: a Cross shape (+) and an Hourglass shape (X). The cross shape was quickly abandoned because the helmet is intended to support lateral interactions with the environment on both sides. To accommodate this functionality, flat mounting surfaces can be integrated on both sides of the helmet for lateral sensor placement. To preserve adequate spacing between frontal and lateral sensing areas, the hourglass (X) arrangement was deemed more appropriate. The resulting geometry prevents overlap between frontal and side sensing regions and supports modular expansion to multi-directional interaction without mechanical interference.

Another decision variable was the relative angle of the helmet's face with respect to the base. To solve this problem, we resorted to simulation. We created 6 hourglass-style helmets, and varied the angle of the face relative to the base, i.e., the pitch of each sensor, in increments of 5° from 10° to 30°, in addition to a 0° option. The custom environment created in the Genesis physics simulator [11] featured a Unitree Go1 robot with the Head-On helmet mounted on the robot's head. Four force-contact sensors were placed on the links of the sensors to measure interactions with the environment. For the interaction object, we created a custom mosaic-pattern wall. This enabled a repeatable, controllable data acquisition process, providing a fair comparison across face angles.

B. Optimization Criteria

In order to conduct our simulations to optimize sensor placement, we trained a Deep Reinforcement Learning based locomotion policy π_l using the Proximal Policy Optimization PPO [12] algorithm. The policy receives observations $o_t = [x_t, c_t]$ from the environment and outputs next joint-angle targets as actions a_t. The first in the observation, $x_t = [\mathbf{v}_{xy}, \omega_z, \mathbf{q}, \dot{\mathbf{q}}, a_{t-1}]$ is the proprioceptive state of the robot, where $\mathbf{v}_{xy} \in \mathbb{R}^2$ is the forward and lateral velocity of the robot in the base frame, $\omega_z \in \mathbb{R}$ is the angular velocity about the upward Z axis, $\mathbf{q} \in \mathbb{R}^{12}$ and $\dot{\mathbf{q}} \in \mathbb{R}^{12}$ are the joint angles and joint velocities and $a_{t-1} \in \mathbb{R}^{12}$ is the previous action taken by the policy. The second term, $c_t = [v_c, h_c]$, contains the high-level command terms, where $v_c \in \mathbb{R}$ is the linear velocity command in the forward direction of the robot and the base height command $h_c \in \mathbb{R}$ is the height command. Motivated by recent advances in Deep Reinforcement Learning DRL [13] and specifically with applications in robotics applications with interactions [14], [15], we employ a structured state space sequence model-based architecture for the policy. The policy is trained following a standard set of rewards where the agent is rewarded based on how well they track the commands given to them [16].

The mosaic-style wall was designed to test each of the criteria and simulate various real-world collision scenarios, allowing us to accurately assess the effectiveness of each helmet under different contact conditions.

A total of $20,000$ trials were conducted. In each trial, the robot maintained a constant height command of $h_c = 0.15$ m until wall contact was established; immediately following contact, the command transitioned to a sinusoidal trajectory. The robot was commanded with a forward velocity of $v_c = 0.5$ m/s to initiate direct contact with the wall. Each interaction lasted $T = 10$ s and simulation data were acquired at a frequency of 200 Hz. To ensure the robustness of the results, the robot's initial heading was randomized within the range $[-20, 20]°$, introducing variability in the approach angle and subsequent contact dynamics.

Following data collection, we performed a comparative analysis of the helmet designs using the following two metrics to rank performance:

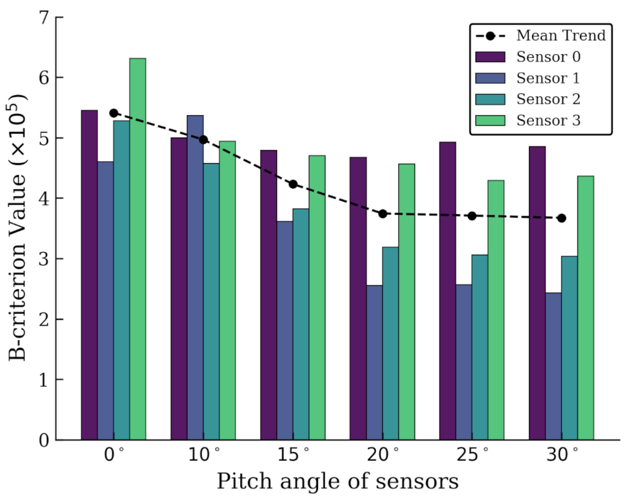

Fig. 3: Helmet shape optimization results for the B-criterion. The helmet with the 0° pitch angle outperformed all other designs in this metric. The designs were ranked for this criterion based on the tallest values in the respective histogram while remaining equally distributed across all four sensors.

1) B-Criterion: The design objective requires a balanced distribution (B) of interactions across the sensor array, minimizing any operational bias toward a specific sensor unit. Let $F_i(t_k;\psi) \in \mathbb{R}^3$ denote the contact force vector recorded by sensor $i \in \{0,1,2,3\}$ at time step t_k for a helmet design parameterized by pitch angle ψ. A force reading is deemed significant if its magnitude exceeds a threshold of $F_{\text{thresh}} = 50$ N. This threshold was selected to isolate purposeful whole-body environmental interactions. For each sensor i, the total sum of significant force readings $N_i(\psi)$ is accumulated across multiple simulation episodes. The candidate designs were ranked to identify the optimal configuration ψ^* that maximizes the mean sensor utilization $\bar{N}(\psi) = \frac{1}{4}\sum_{i=0}^{3} N_i(\psi)$ while minimizing the cross-sensor discrepancy, calculated as the standard deviation $\sigma_N(\psi)$. As shown by the individual sensor distributions in Fig. 3, the 0° pitch angle configuration simultaneously maximizes the total deliberate contact frequency $\bar{N}(\psi)$ and minimizes the directional bias $\sigma_N(\psi)$.

2) SNR-Criterion: For each sensor, the planar shear forces must exhibit a zero-mean normal distribution with a maximized Signal-to-Noise Ratio (SNR). Specifically, an unbiased contact surface, meaning there is no preferential sliding direction, requires the planar force vector $\mathbf{F_{xy}} = [F_x, F_y]^T$ to follow $\mathbf{F_{xy}} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma_{xy}})$, where $\mathbf{0}$ is the zero mean vector and $\mathbf{\Sigma_{xy}}$ is the covariance matrix of the forces. To quantify the quality of this interaction, we define the planar SNR as the ratio of the force magnitude to the spatial variance:

$$SNR_{xy} = \frac{\|\mathbf{F_{xy}}\|}{\sqrt{\text{tr}(\mathbf{\Sigma_{xy}})}} = \sqrt{\frac{F_x^2 + F_y^2}{\lambda_1 + \lambda_2}} \tag{1}$$

Here, $\|\mathbf{F_{xy}}\|$ is the Euclidean norm of the force vector, $\text{tr}(\mathbf{\Sigma_{xy}})$ is the trace of the covariance matrix, and λ_1, λ_2

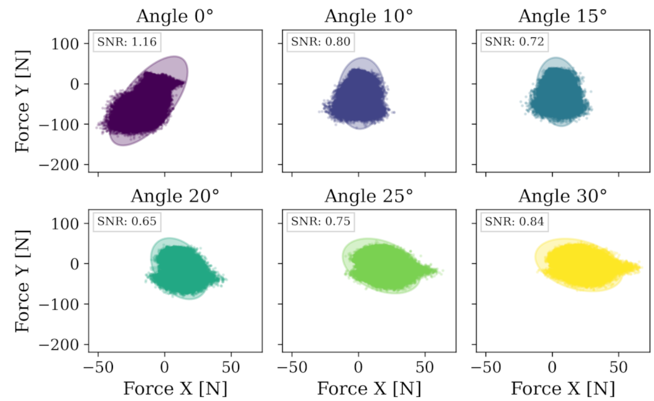

Fig. 4: Helmet shape optimization results for the SNR-criterion for Sensor 0. The helmet with 0° pitch angle outperformed all other designs. The designs were ranked for this criterion based on the largest SNR ratio.

are its eigenvalues. Geometrically, these eigenvalues dictate the variance along the principal axes of the confidence ellipse that bounds 99.7% of the force readings (the 3σ threshold). To this end, a larger SNR_{xy} indicates a stronger, more deliberate contact signal relative to the multi-directional noise.

For brevity, Fig. 4 visualizes the distribution results exclusively for the top-right sensor (Sensor 0), while the comprehensive SNR values for all sensors across varying helmet face angles are detailed in Table I. The data demonstrate that the 0° helmet configuration yields the highest SNR across almost all sensors, placing a close second only for Sensor 3. Consequently, the 0° design emerges as the optimal configuration in this comparative analysis.

TABLE I: Signal-to-Noise Ratio (SNR) by Sensor and Pitch Angle

Pitch Angle	Sensor 0	Sensor 1	Sensor 2	Sensor 3
0°	**1.16**	**1.16**	**1.10**	0.96
10°	0.80	0.65	0.46	0.70
15°	0.72	0.61	0.44	0.66
20°	0.65	0.80	0.60	0.71
25°	0.75	0.98	0.81	0.90
30°	0.84	1.03	0.87	**1.01**

IV. Tactile Sensor Calibration

From our experience with the AnySkin sensors, we found that significant drift in the sensor's raw measurements is primarily due to permanent skin deformation following continuous loading and unloading. Previous work employed a static baseline calculation at the beginning of an experiment. This is done by measuring the average of N_{bl} samples $\bar{B} = \frac{1}{N_{bl}}\sum_{j=0}^{N_{bl}} B_j$ [8], [9]. Then each subsequent measurement B_t is debiased by subtracting this baseline to produce the corrected measurement $B'_t = B_t - \bar{B}$. However, this baseline approach presents critical limitations when applied to dynamic, floating-base systems. First, it only compensates for geomagnetic interference under nearly static conditions. During active locomotion, the rapid orientation changes of the robot's base, which routinely reach angular

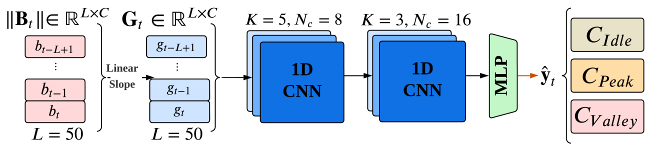

Fig. 5: The proposed cuasal 1D-CNN for peak and valley detection. The input to the model is a sequence, with length L, of linear slopes $\mathbf{G}_t$ extracted from the magnitude of the magnetic flux of all sensors. The model predicts the sequence's class as Idle, Peak, or Valley.

velocities of up to 180 $°/s$ [17]–[19], induce significant drift in the raw magnetometer data, rendering the initial static baseline obsolete within seconds. Second, prior literature primarily evaluates these sensors on fixed-base manipulators, where orientation-induced magnetic shifts can be analytically predicted and mitigated using known forward kinematics. This predictive compensation is infeasible for autonomous mobile robots operating in unstructured environments, as their continuous trajectory and absolute global orientation cannot be exhaustively anticipated ahead of time.

To address these dynamic challenges, we propose a novel two-step calibration framework. The proposed methodology is robot- and sensor-agnostic and thus represents a generalizable calibration methodology for magnetometer-based tactile sensors deployed on floating-base platforms. The central piece of this process is a robust contact estimation model, detailed in Algorithm 1. Accurate contact state detection is critical: by reliably isolating no-contact phases, the system can dynamically compute a new baseline to eliminate accumulated drift and hysteresis. This continuous cycle of state estimation and dynamic re-baselining ensures that the magnetic flux measurements remain within a valid operational domain, enabling a consistent and reliable mapping to environmental contact forces.

A. Real-Time Contact Estimation via Causal-1D CNN

A fundamental challenge in soft tactile and magnetic skin sensors is the inherent drift and hysteresis present in the raw amplitude signals, which render static thresholding unreliable for continuous contact estimation. To achieve robust, real-time calibration and contact detection, we propose a dynamic, event-driven architecture utilizing a strictly Causal 1D Convolutional Neural Network (CNN) as the backbone for sequence classification, shown in Fig. 5. The algorithm is presented in Algorithm 1.

A fundamental advantage of this approach is that, rather than relying on static thresholds applied to raw magnetic flux amplitudes, the model evaluates the temporal dynamics of the signal's trend. Let $\mathbf{B}_f(t) \in \mathbb{R}^{C\times 3}$ denote the filtered magnetic flux derived from the raw readings $\mathbf{B}(t)$ for C active sensor channels at timestep t. In our setup we use 4 AnySkin sensors and each one has 5 magnetometers thus the number of channels is $C = 20$. To capture these dynamics, the algorithm maintains a primary sliding window of 50 samples ($L = 50$), extracting the instantaneous trend, g_t, via linear least-squares polynomial fitting. A secondary buffer then aggregates these temporal gradients into a sequential feature vector, defined as $\mathbf{G}_t = [g_{t-L+1}, \dots, g_t] \in \mathbb{R}^{L\times C}$,

Algorithm 1 Real-Time Causal CNN Contact Estimation & Calibration

Require: Active channels $C \in \mathbb{N}_{>0}$, Raw sensor measurement $B_t \in \mathbb{R}^{C\times 3}$, Initial baseline $\bar{B} \in \mathbb{R}^{C\times 3}$, Per-channel offset vector $O \in \mathbb{R}^{C\times 3}$, Queue Length L.
Ensure: Updated contact state $S \in \{0,1\}^C$, Updated channel offsets $O \in \mathbb{R}^{C\times 3}$

Initialize: $W_b \leftarrow \emptyset$, $G \leftarrow \emptyset$ ▷ Fixed-size queues
Initialize: $S \leftarrow \{0\}^C$, $S_{prev} \leftarrow \{0\}^C$
procedure UPDATE(B_t) ▷ New measurement
  $B'_t \leftarrow B_t - \bar{B}$ ▷ Subtract initial baseline
  $B_{mean} \leftarrow \text{mean}(B'_t, axis = 0)$ ▷ Per channel mean
  $B_{rebased} \leftarrow B_{mean} - O$ ▷ Subtract channel offsets
  $B_f \leftarrow \text{LowPassFilter}(\text{MedianFilter}(B_{rebased}))$
  W_b.Push(norm(B_f))
  if length(W_b) $= L$ **then** $g_t \leftarrow \emptyset$
    for $i = 1$ **to** C **do**
      $g_t[i] \leftarrow \text{LeastSquaresSlope}(W_b[i])$
      g_t.Push($g_t[i]$) ▷ Per channel slope
    G.Push(g_t)
    if length(G) $= L$ **then**
      $\mathbf{y}_t \leftarrow f_\theta(G)$ ▷ Causal CNN fwd pass
      $\hat{\mathbf{c}}_t \leftarrow \text{argmax}(\mathbf{y}_t)$
      $S_{prev} \leftarrow S$ ▷ Store previous contact state
      for $i = 1$ **to** C **do**
        if $\hat{c}_i = 1$ **then** ▷ Peak detected
          $S_i \leftarrow$ True ▷ Set contact
        else if $\hat{c}_i = 2$ **then** ▷ Valley detected
          $S_i \leftarrow$ False ▷ Unset contact
        if $S_{prev,i}$ **and not** S_i **then**
          $O[i] = \text{RebaselinePatch}(i)$ ▷ Update channel offsets on falling edge

which provides the necessary temporal context for robust contact estimation.

Another key observation, related to the chosen network architecture, is that standard centered convolutions would produce data leakage from future time-steps, rendering the method unusable for online robotic setups. To satisfy strict real-time constraints, the CNN backbone is constructed with causality into its structure, enforced through asymmetric left-padding. For a temporal convolution with kernel size K and dilation factor d, the input sequence is padded with $P = (K-1)\times d$ zeros on the left side (past) of the temporal dimension.

The network architecture consists of an empirical normalization layer utilizing a global scalar mean and variance, followed by two causal convolutional blocks. The first block consists of a causal convolutional layer with kernel size $K = 5$, channel dimension $N_c = 8$, stride $S = 2$, and a Gaussian Error Linear Unit (GELU) activation. The second block contains a causal convolutional layer with kernel size $K = 3$, channel dimension $N_c = 16$, stride $S = 2$, and a GELU activation. This architecture extracts hierarchical temporal features without the spatial misalignment risks

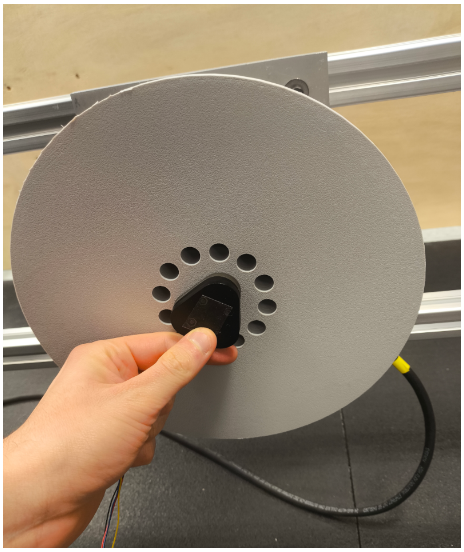

Fig. 6: Calibration setup for the AnySkin tactile sensors. A single AnySkin sensor is pressed against a custom 3D-printed sensing plate rigidly mounted to an ATI Axia130 M125 6-axis force-torque sensor, which provides ground-truth measurements. During data collection, the dynamic re-baselining strategy is actively applied to mitigate sensor drift, ensuring the raw magnetic flux readings remain within a valid domain for accurate force mapping.

associated with standard max-pooling in causal sequences.

The network is optimized to classify sequences into 3 possible classes: Contact Onset ($C_{peak} = 1$), Contact Release ($C_{valley} = 2$), and Idle ($C_{idle} = 0$). Let f_θ represent the parameterized network, mapping the spatial-temporal input to class logits:

$$\hat{\mathbf{c}}_t = f_\theta(\mathbf{G}_t) \in \mathbb{R}^3 \tag{2}$$

We used the Categorical Cross-Entropy loss function to optimize the weights of the neural network:

$$\mathcal{L}_{CE} = -\sum_{t=0}^{L} y_t \log(\hat{y}_t) = -\sum_{t=0}^{L} y_t \log(f_\theta(\mathbf{G}_t)), \tag{3}$$

where $y_t \in 0, 1, 2$ is the ground truth label, $\hat{y}_t \in 0, 1, 2$ is the classification prediction from the model at time t computed as $\hat{y}_t = \text{softmax}(\hat{\mathbf{c}}_t)$. By mapping the temporal evolution of the signal gradients to discrete latent events, the CNN effectively serves as the trigger mechanism for a robust Boolean state machine. An onset prediction ($\hat{y}_t = C_{peak}$) transitions the sensor state to active contact, while a release prediction ($\hat{y}_t = C_{valley}$) resets the state. The idle class (C_{idle}) acts as a zero-order hold, inherently bypassing the need for continuous amplitude thresholding and rendering the system highly resilient to sensor hysteresis and baseline drift during prolonged physical interactions. Training of the CNN model takes about 10 seconds and can be rapidly trained with more data if available and necessary.

To train the sequence classification model, we compiled a comprehensive dataset of continuous multi-channel magnetic sensor readings by capturing both human-guided and autonomous robot interactions with everyday objects, such as office furniture and a suitcase. This diverse corpus captures a wide range of contact morphologies and dynamic loading profiles, ensuring the model generalizes effectively to real-world physical engagements.

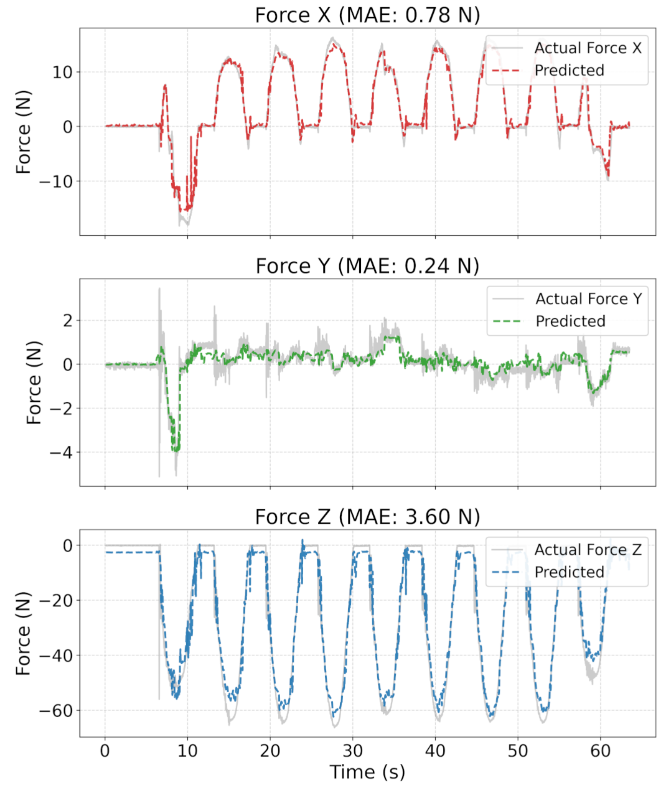

Fig. 7: Performance of the lightweight LSTM-MLP neural network mapping magnetic flux magnitudes to 3D contact forces. When evaluated against ground-truth force sensor data, this computationally efficient architecture achieves a low Mean Absolute Error (MAE) across all axes, demonstrating sufficient fidelity for real-time force estimation.

Given the raw magnetic sensor readings from the dataset, we extract ground truth contact transitions using a prominence-based peak detection algorithm [20], yielding discrete sets of peak ($\mathcal{I}_{\text{peak}}$) and valley ($\mathcal{I}_{\text{valley}}$) indices. Crucially, the ground truth generation algorithm cannot be used in real time as it utilizes future information to detect peaks and valleys. Finally, to prevent class imbalance and establish a baseline, idle state indices ($\mathcal{I}_{\text{idle}}$) are randomly sampled from temporal regions strictly outside the active event windows.

B. Force Mapping & Calibration

Motivated by prior tactile calibration methodologies [9], establishing a reliable force mapping requires robust ground truth measurements. To achieve this, we utilized an ATI Axia130 M125 6-axis force-torque sensor to record reference contact forces (Fig. 6). The training dataset was generated by pressing a single AnySkin sensor against the force plate while actively applying the dynamic re-baselining strategy detailed in Algorithm 1.

To model this relationship, we trained a lightweight architecture comprising a Long Short-Term Memory (LSTM) network and a Multi-Layer Perceptron (MLP). This network,

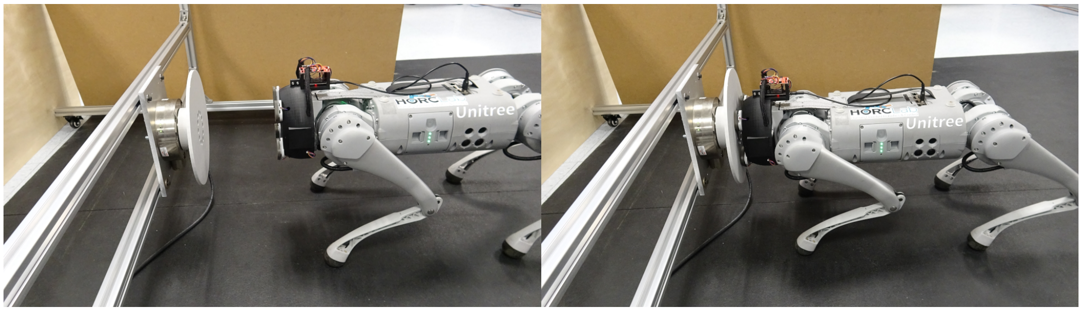

Fig. 8: Experimental sequence of the controlled pushing evaluation. Left: The quadruped robot in its initial pre-contact stance, positioned opposite the ATI force-torque sensor plate. Right: Active contact phase, illustrating the sensorized helmet engaging the force plate during dynamic fore-aft loading.

f_ϕ, maps the magnetic flux vector $\mathbf{B}_t$ directly to the estimated 3D force vector $\hat{\mathbf{F}}_t$:

$$\hat{\mathbf{F}}_t = f_\phi(\mathbf{B}_t) \in \mathbb{R}^3 \tag{4}$$

As illustrated in Fig. 7, the trained network achieves a Mean Absolute Error (MAE) of 3.6 N in the Z-direction that is the most critical for our application. While more complex architectures could potentially yield higher mapping fidelity [8], [9], this lightweight model provides a computationally efficient and sufficiently accurate baseline for real-time robotic deployment.

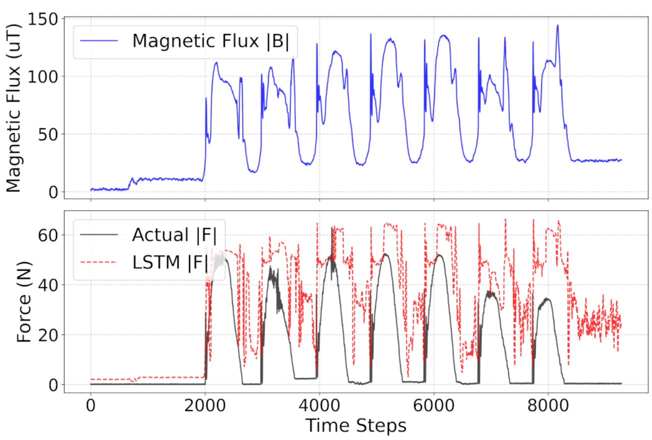

Fig. 9: System performance without dynamic re-baselining. Top: Raw magnetic flux magnitude for Sensor 2, exhibiting severe drift accumulation over successive cyclic interactions. Bottom: The resulting force prediction from the LSTM-MLP model compared to ground-truth force measurements. The unmitigated sensor drift fundamentally corrupts the network input, leading to a rapid and compounding degradation in force mapping accuracy.

V. Experimental Results

The primary objective of our experimental evaluation is to demonstrate the robustness and repeatability of the proposed tactile sensing architecture under dynamic loading conditions. As illustrated in Fig. 8, the experimental testbed comprises the quadruped robot, the sensorized helmet, and an ATI force-torque sensor used to establish ground-truth physical interactions. To rigorously validate the combined contact-detection and dynamic re-baselining strategy, we conducted a series of controlled push experiments. The system's software pipeline, implemented via custom ROS2 packages and PyTorch [21], facilitates seamless, low-latency data streaming at 150 Hz from the robot's onboard compute module to an external unit, ensuring consistent real-time performance throughout the validation trials.

A. Controlled Push Experiment

To evaluate the efficacy of the proposed contact estimation and dynamic re-baselining framework, we designed a controlled kinematic experiment. A trajectory planner was implemented to stream target positions to the robot's low-level Proportional-Derivative (PD) controllers. Starting from a nominal standing posture, synchronous sinusoidal trajectories were commanded to the four sagittal hip actuators. This control strategy induced a continuous fore-aft translation of the robot's floating base while maintaining fixed foot-ground contact constraints throughout the duration of the trial.

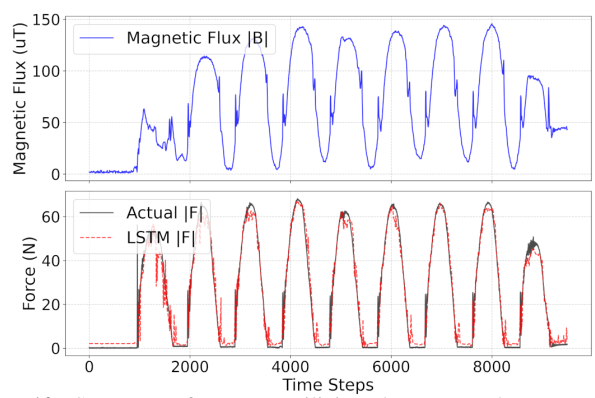

Fig. 10: System performance utilizing the proposed contact estimation and dynamic re-baselining framework. Top: Raw magnetic flux magnitude for Sensor 2, demonstrating effective drift mitigation through state-dependent zeroing. Bottom: The LSTM-MLP force prediction compared to ground-truth force measurements. By continuously anchoring the raw signal to its valid operational domain, the proposed method sustains high-fidelity force estimation across extended, repeated interactions.

To demonstrate the necessity and efficacy of our approach, we compared the system's performance with and without the dynamic re-baselining framework. As illustrated in Fig. 9, operating the robot without contact estimation results in severe, unmitigated drift in the raw magnetic flux measurements. This rapid accumulation of systemic noise corrupts the inputs to the LSTM-MLP network, causing the force predictions to severely degrade and deviate from reality within a few actuation cycles. Conversely, Fig. 10

highlights the system's robust response when the proposed contact estimation and dynamic re-baselining algorithm is active. The continuous, state-dependent zeroing effectively eliminates drift, ensuring that the raw magnetic signals remain strictly within their calibrated domain. Consequently, the network's force predictions maintain high fidelity and closely track the ground-truth ATI sensor data with an MAE of $8\ N$, even throughout continuous and repeated physical interactions.

VI. Conclusion

In this study, we developed a novel tactile sensing helmet for the Unitree Go1 quadruped robot, enabling it to interact with its environment in previously unexplored ways. More broadly, this work serves as a foundation for future research aimed at enhancing the environmental and interactive capabilities of quadruped robots. Additionally, we presented a novel contact estimation algorithm for magnetometer-based tactile sensors. This algorithm effectively manages the complexities arising from a dynamic, moving base, thereby overcoming significant limitations of prior calibration methods. Our experimental results demonstrate that the proposed method effectively mitigates inherent sensor drift, a persistent challenge that has historically hindered the application of these tactile sensors [6]–[9].

While the proposed tactile sensing framework demonstrates significant robustness in dynamic settings, certain limitations remain. Currently, the contact estimation algorithm occasionally yields false negatives during highly transient interactions. Future research will mitigate this by leveraging expanded, more diverse contact datasets and exploring more sophisticated sequence-modeling architectures to improve detection accuracy. Additionally, translating these continuous, high-fidelity force estimates into actionable feedback for advanced locomotion frameworks, specifically Model Predictive Control (MPC) and Deep Reinforcement Learning (DRL) policies, represents a critical next step for achieving closed-loop, whole-body tactile control.

Despite these current constraints, this work establishes a fundamental tactile capability for floating-base platforms operating outside structured laboratory environments. By providing reliable, drift-free force feedback during dynamic locomotion, this sensorized helmet and the novel dynamic re-baselining methodology that supports it lay the groundwork for safe, robust, and deliberate physical Human-Robot Interaction (pHRI). Ultimately, equipping quadrupedal robots with such a resilient sense of touch will be instrumental in accelerating their seamless and safe integration into complex, human-centric environments.

References

[1] Y. Fan, Z. Pei, C. Wang, M. Li, Z. Tang, and Q. Liu, "A review of quadruped robots: Structure, control, and autonomous motion," *Advanced Intelligent Systems*, vol. 6, no. 6, p. 2300783, 2024. [Online]. Available: https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/aisy.202300783

[2] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, p. eabc5986, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.abc5986

[3] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.abk2822

[4] T. Van Hauwermeiren, A. Coene, and G. Crevecoeur, "Tactile force sensing for admittance control on a quadruped robot," *Machines*, vol. 13, no. 5, 2025.

[5] M. Meribout, N. Abule Takele, O. Derege, N. Rifiki, M. El Khalil, V. Tiwari, and J. Zhong, "Tactile sensors: A review," *Measurement*, vol. 238, p. 115332, 2024.

[6] Z. Zhan, Y. Yang, W. Zuo, M. Xie, and M. Ning, "Recent advances and challenges of tactile sensing for robotics: from fundamentals to applications," *Materials Today Physics*, vol. 54, p. 101740, 2025.

[7] J. Xi, H. Yang, X. Li, R. Wei, T. Zhang, L. Dong, Z. Yang, Z. Yuan, J. Sun, and Q. Hua, "Recent advances in tactile sensory systems: Mechanisms, fabrication, and applications," *Nanomaterials*, vol. 14, no. 5, 2024.

[8] R. M. Bhirangi, V. Pattabiraman, E. Erciyes, Y. Cao, T. L. Hellebrekers, and L. Pinto, "Anyskin: Plug-and-play skin sensing for robotic touch," *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 16 563–16 570, 2024.

[9] R. M. Bhirangi, T. L. Hellebrekers, C. Majidi, and A. Gupta, "Reskin: versatile, replaceable, lasting tactile skins," *CoRR*, vol. abs/2111.00071, 2021. [Online]. Available: https://arxiv.org/abs/2111.00071

[10] Espressif Systems, "ESP32 WROOM-32 Datasheet," https://www.espressif.com/en/products/socs/esp32, 2024, accessed: 2024-05-20.

[11] Genesis Authors, "Genesis: A generative and universal physics engine for robotics and beyond," https://github.com/Genesis-Embodied-AI/Genesis, Dec 2024, gitHub repository.

[12] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[13] C. Lu, Y. Schroecker, A. Gu, E. Parisotto, J. Foerster, S. Singh, and F. Behbahani, "Structured state space models for in-context reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 36, 2023.

[14] S. E. Bachoumas and P. Artemiadis, "Learning adversarial policies for swarm leader identification using a probing agent," in *2025 International conference on robotics and automation (ICRA)*. IEEE, 2025.

[15] S. E. Bachoumas and P. Artemiadis, "On swarm leader identification using probing policies," 2025. [Online]. Available: https://arxiv.org/abs/2512.18146

[16] S. Jeon, M. Jung, S. Choi, B. Kim, and J. Hwangbo, "Learning whole-body manipulation for quadrupedal robot," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 699–706, 2024.

[17] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, and A. D. Prete, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2024.

[18] B. Katz, J. D. Carlo, and S. Kim, "Mini cheetah: A platform for pushing the limits of dynamic quadruped control," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6295–6301.

[19] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the mit cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1–9.

[20] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.

[21] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," 2019. [Online]. Available: https://arxiv.org/abs/1912.01703