Integrating Reconfigurable Foot Design, Multi-modal Contact Sensing, and Terrain Classification for Bipedal Locomotion

Ted Tyler [*], **Vaibhav Malhotra** [*], **Adam Montague** [*],
Zhigen Zhao [*], **Frank L. Hammond III** [*], and **Ye Zhao** [*]

[*] *George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30309 USA (e-mail: ttyler31, vaibhavmalhotra99, adam.montague, zhigen.zhao@gatech.edu and frank.hammond, ye.zhao@me.gatech.edu).*

Abstract: The ability of bipedal robots to adapt to diverse and unstructured terrain conditions is crucial for their deployment in real-world environments. To this end, we present a novel, bio-inspired robot foot design with stabilizing tarsal segments and a multifarious sensor suite involving acoustic, capacitive, tactile, temperature, and acceleration sensors. A real-time signal processing and terrain classification system is developed and evaluated. The sensed terrain information is used to control actuated segments of the foot, leading to improved ground contact and stability. The proposed framework highlights the potential of the sensor-integrated adaptive foot for intelligent and adaptive locomotion.

Keywords: Mechatronic Systems; Robotics; Modeling, Identification and Signal Processing;

1. INTRODUCTION

Legged robots are increasingly required to traverse a variety of environments and terrains to accomplish their tasks (Shamsah et al. (2023)). For example, agricultural robots for fruit harvesting need to walk on grass, dirt, mud, snow, and mulch; disaster recovery robots could navigate through building rubble, ruined streets, and collapsed houses; and home-assistance robots may traverse rugs, carpets, and a variety of tile and hardwood floors. For each type of terrain, the robot foot should be robust enough to enable stable ground support. In addition, sufficient sensory input and processing capability are necessary to identify the terrains in order to adjust the control strategy of the robot. This paper proposes a versatile foot design for the Cassie robot, integrated with a multi-modal sensor suite (see Fig. 1), and employs machine learning algorithms to classify terrains based on the sensor data.

Terrain-adaptive foot design is essential for a bipedal robot's ability to traverse harsh terrains. The default foot of a Cassie robot is narrow and prone to sinking on deformable terrains. In contrast, a rigid wide foot design does not make full contact with the ground if the terrain is not flat. Therefore, we propose a nature-inspired foot design with tarsal segments that can be actuated according to the terrain types, which adapts to terrain surfaces with soft and uneven properties.

The robust classification of the terrain relies on multiple sources of sensory input. The features extracted from the sensor data are input into a machine learning based terrain classification system. The signal acquisition and terrain classification framework is closely related to the previous work in Guo et al. (2020), but with a newly designed sensor suite involving temperature, acoustic, acceleration, capacitive, tactile measurements. Additionally, new fea-
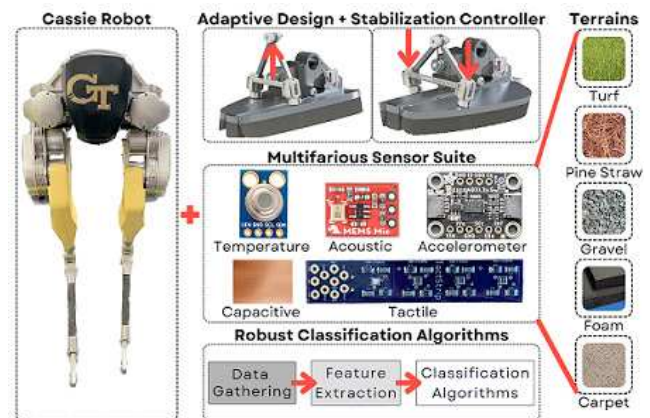

Fig. 1. Illustration of a reconfigurable Cassie foot design with stabilizing tarsal segments and multi-modal sensing suite (i.e., multiple sensors for various features).

ture extraction methods and deep learning algorithms are explored in this study.

The contributions for this paper lie in: (i) a nature-inspired foot design with adaptive foot stabilization over uneven terrains, complete with a high-level controller; (ii) the design of a tactile sensor with commercially available components, as well as the integration of the multifarious sensor suite and associated electronics; (iii) a terrain identification system capable of real-time operation on a Cassie robot, and an evaluation of the performances of various supervised learning algorithms for terrain classification.

2. RELATED WORK

Adaptive foot design: Previous attempts have been made to design an adaptive foot with sensor integration to enhance robot locomotion. The works in Catalano et al.

(2021) and Piazza et al. (2016) propose articulated soft foot designs that are compliant with unstructured terrains to maximize contact. The work in Fondahl et al. (2012) develops an actuated multi-contact foot with pressure sensor arrays integrated.

Sensors for terrain classification: A combination of inertial measurement units (IMUs), force, and torque sensors are commonly used to gather terrain data. Examples of this include Kolvenbach et al. (2019), Wang et al. (2020), Bosworth et al. (2016), Venâncio et al. (2020), Walas (2015), and Walas et al. (2016). Visual data is also a popular input: Walas (2015) combines force and torque data with visual and depth data, Halatci et al. (2008) uses a combination of images and wheel vibration in a rover, Khan et al. (2011) and Khan et al. (2012) both use a single robot-mounted camera, and Kurobe et al. (2021) uses camera and microphone input to identify terrain characteristics. Tactile sensors, which mimic a human's ability to sense through touch, are also beneficial for terrain classification. Guo et al. (2020) proposes to use a variety of embedded sensors including tactile, acoustic, capacitive, temperature, and acceleration sensors, and shows that integrating multiple sensor modalities in general results in more effective terrain identification.

Machine learning algorithms for terrain classification: Support vector machine is used by Kolvenbach et al. (2019), Walas (2015), Walas et al. (2016), Halatci et al. (2008), and Wu et al. (2016), while Khan et al. (2011), Khan et al. (2012), and Zhang et al. (2016) all use the random forest algorithm. Ding et al. (2022) experiments with four traditional algorithms: support vector machine, logistic regression, k-nearest neighbor, and gaussian discriminant analysis. Recent trends indicate the potential of deep learning algorithms for terrain identification due to their enhanced feature learning capabilities compared to traditional machine learning methods as well as their ability to generalize over large quantities of data. Artificial neural networks (ANNs) are used in Giguere and Dudek (2011); Convolutional neural networks (CNNs) are used in Venâncio et al. (2020), Kurobe et al. (2021), and Vulpi et al. (2021); Recurrent neural networks (RNNs) are used in Wang et al. (2020) and Vulpi et al. (2021); and concepts from transformers are used in Bednarek et al. (2022) and Sójka et al. (2023).

3. METHODS

This section is organized in the following manner: the mechanical design and stability control of the adaptive foot is discussed in Sec. 3.1; the electrical design and sensor integration are discussed in Sec. 3.2; the signal processing and terrain classification framework is discussed in Sec. 3.3.

3.1 Mechanical Design

Two primary objectives are considered for the design of the foot. First, the foot must incorporate new tarsal segments that can adaptively stabilize the robot on rough terrain. Second, the foot needs to be lightweight to minimize deterioration of the Cassie robot's joints.

The design shown in Figs. 1 and 2 fulfills the first objective by efficiently transferring rotational force from a servo into two tarsal segments on either side of the foot that act on

the ground. The servo located at the center of the foot rotates to a selected angle between 0 and 360° based on the feedback from the sensors and the classification algorithm. As the servo rotates, the reinforced string attached to the servo's output shaft winds up, pulling the joint above the center of the foot down. This pulls on the linkages which directly push the tarsal segments down. The tarsal segments, which are attached to the central part of the foot using hinges, rotate to a maximum angle of fifteen degrees below the horizontal. When the foot is raised off the ground, the servo unwinds and the spring connecting the two tarsal segments passively retracts the tarsal segments to ten degrees above the horizontal. As such, the force from the servo acts through the linkages to pull directly on the tarsal segments. While it does have to overcome the force of the spring, it is a relatively weak spring that serves to barely lift the tarsal segments off the ground, preserving as much force from the servo for the tarsal segments as possible.

The servo provides a maximum torque of 22.8 kg-cm, which is converted through the linkages in our design into a stabilizing force of up to 447 N split between the two tarsal segments. If the robot is perfectly balanced and the tarsal segments are both flat on the ground, this force will be split evenly between the segments; as the robot falls to one side, the force will shift to the tarsal segment on that side to help the robot stay upright.

To actuate the tarsal segments, a bang-bang controller, running at 45 Hz with the sensors, uses feedback from the classification algorithm, represented by the adaptive design portion of Fig. 1. The servo operates the tarsal segments if the algorithm predicts a terrain that is potentially destabilizing: poppy seeds, gravel, pine straw, turf, a foam block, or carpet. As the tarsal segments descend, they increase the overall contact area of the foot by 292% to distribute weight and stabilize the robot. As loose terrains shift, the tarsal segments can descend further to maintain contact with the ground. If a stable terrain is classified, the tarsal segment remains in its upright, passive position and only the central part of the foot is in contact. The current controller is limited to being either fully-actuated or passive, meaning that the angle cannot be fine-tuned based on the terrain classification. The Cassie robot takes a step every 0.4 seconds. Thus, running the controller at 45 Hz means there is ample time to react to any new input from the classification algorithm between steps.

To achieve the goal of a lightweight foot, the physical components are fabricated using a variety of materials and methods. The tarsal segments, linkages, and actuator
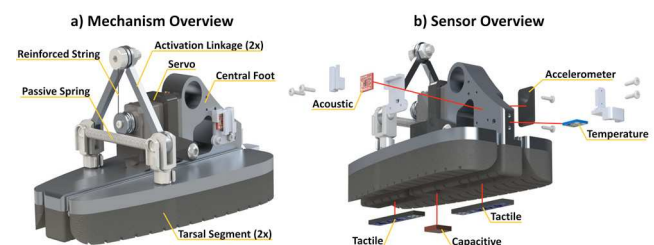

Fig. 2. A labeled rendering of the foot that shows the placement of sensors and mechanical components. The red lines show the sensor placements on the foot.

mount are all 3D printed using a nylon and carbon fiber filament, while the central component of the foot is milled using Aluminum 6061-T6. Other, less structurally essential components, like sensor mounts, are 3D printed using ABS filament. The final mass of all mechanical components of the foot is 0.64 kg, while the total mass including electronics is 1.05 kg. For reference, the original Cassie robot's foot is 0.21 kg.

3.2 Electrical Design

Five sensors are used on the foot: capacitive, acoustic, tactile, temperature, and accelerometer. These sensors are shown on the foot in Fig. 2. The capacitive sensor is a copper plate connected to the microcontroller with a wire while the acoustic sensor (SPH8878LR5H-1, SparkFun), accelerometer (ADXL343, Adafruit), and temperature sensor (GY-906 MLX90614ESF, HiLetgo) are commercially available breakout boards. The new tactile sensors, like the one in Fig. 3, are based on the design by Tenzer et al. (2014) that uses barometric pressure sensors to register forces. We follow a similar fabrication process to the one used by Kõiva et al. (2020) to construct the tactile sensors. However, two key differences from these papers are that we use Bosch BMP390 sensors (Bosch Sensortec), which communicate with the microcontroller through an I2C Multiplexer (TCA9548A, Adafruit) capable of coordinating with up to eight sensors with the same I2C address at once. To take full advantage of this multiplexer, two tactile sensors are used on the foot and each tactile sensor has four Bosch sensors.

A Raspberry Pi 4 microprocessor is used to provide the processing power to perform data acquisition and terrain classification in real time. However, the Pi lacks analog to digital converters, making it incompatible with the requirements for the acoustic and capacitive sensors. Therefore, a Teensy 3.6 microcontroller is used with the Raspberry Pi. The Teensy gathers the data from the capacitive and acoustic sensors at 18 kHz and performs feature extraction. The Raspberry Pi reads and processes data from the other sensors at 45 Hz, and requests features extracted on the Teensy at the same frequency.

Each sensor needs to be compatible with the Raspberry Pi and the Teensy, which have 3.3 v and 5.0 v supply voltages. Each sensor uses 3.3 v, excluding the temperature sensor, which operates at 5.0 v. Individually, the sensors need to have appropriate ranges to be able to respond to the different stimuli that are likely to occur on a given terrain. The temperature sensor has a range of $-40°$ C - 380° C, which more than covers any likely terrain temperature. The new pressure sensor, accelerometer, and acoustic sensor, which replace their obsolete counterparts, all outperform the sensors used by Guo et al. (2020): a
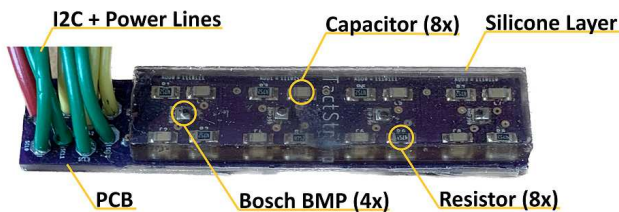
pressure range of 30 kPa - 125 kPa vs 50 kPa - 115 kPa, acceleration resolution options of $\pm$ 2 g /$\pm$ 4 g /$\pm$ 8 g /$\pm$ 16 g vs $\pm$ 2 g /$\pm$ 4 g /$\pm$ 8 g, and a frequency range of 7 Hz - 36 kHz vs 100 Hz - 15 kHz.

3.3 Terrain Classification Algorithms

Preprocessing: The raw data is preprocessed and segmented so that the appropriate features can be extracted for use in classification. Before preprocessing, the raw tactile data may experience linear drift over time. Additionally, an offset is at times introduced to the pressure sensors during the manufacturing processes needed to make the tactile sensors. To remove the drift and keep the input from each pressure sensor at a constant baseline, a detrending function is applied to the input of each sensor. This function performs a linear least-squares fit to the data and subtracts the solution from the data to remove the linear drift. The detrended tactile data is used to segment the rest of the data, i.e., split the data every time a footstep is taken so that statistical features can be extracted around the footsteps, and the unnecessary data between footsteps can be ignored. This is accomplished by identifying maxima and minima in the tactile data that corresponds to the foot in contact with the terrain. [1]

Feature Extraction: Feature extraction is applied to the segmented data so that each footstep is represented by a single set of variables for the machine learning algorithms. Using extracted features instead of raw data reduces the dimensionality of the data and facilitates the classification process. In total, 100 features are extracted from the five sensors using the equations below, where $\mathbf{x} = \{x_0, x_1, \ldots, x_n\}$, and $x_i \in \mathbb{R}$ denotes the i^{th} scalar in a time series of sensor data with length n, and ZCR stands for the Zero-crossing Rate:

Sum: $f_{\text{sum}}(\mathbf{x}) = \sum_{i=0}^{n} x_i$

Maximum: $f_{\text{max}}(\mathbf{x}) = \max(x_1, x_2, x_3, ..., x_n)$

Minimum: $f_{\text{min}}(\mathbf{x}) = \min(x_1, x_2, x_3, ..., x_n)$

Mean: $\mu(\mathbf{x}) = \frac{1}{n} \sum_{i=0}^{n} x_i$

Variance: $\sigma^2(\mathbf{x}) = \frac{1}{n} \sum_{i=0}^{n} (x_i - \mu)^2$

Skewness: $\tilde{\mu}^3(\mathbf{x}) = \frac{1}{(n-1)^3} \sum_{i=0}^{n} (x_i - \mu)^3$

Kurtosis: $f_{\text{kurt}}(\mathbf{x}) = \frac{1}{n(\sigma^2)^2} \sum_{i=0}^{n} (x_i - \mu)^4$

Sign: $f_{\text{sign}}(x_i) = \begin{cases} 0 & x_i < 0 \\ 1 & x_i \geq 0 \end{cases}$

ZCR: $f_{\text{zcr}}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^{n} |f_{\text{sign}}(x_i) - f_{\text{sign}}(x_{i-1})|$

80% Rise Time: $f_{80\%}(\mathbf{x}) = i$ for $x_i = 0.8(f_{\text{max}}(\mathbf{x}))$

Fast Fourier Transform: $f_{\text{fft}}(\mathbf{x}) = FFT(\mathbf{x})$

The features extracted from each sensor are as follows: *Accelerometer:* the maximum, minimum, mean, variance, sum, and ZCR are calculated for each axis (X, Y, and Z) on the three-axis accelerometer for eighteen features.

Fig. 3. Closeup of a completed tactile sensor, with individual components labeled.

---

[1] See https://github.com/ATMontague/Cassie-Robot-Foot-Design for additional information on the electrical design and preprocessing.

Acoustic: the ZCR and Fast Fourier Transform (FFT) functions are calculated. The FFT output is averaged over nine spectral bands for ten features.

Capacitive: the mean and variance are calculated, as well as the same nine spectral band averages that are calculated for the acoustic sensor for 11 features.

Tactile: the maximum value, minimum value, and 80% rise time values are calculated for the data summed at each time step across the eight barometric pressure sensors. We also calculate the maximum, minimum, mean, variance, skewness, and kurtosis values for each individual sensor, as well as each individual sensor's value when the summed value reaches 80% rise time for 59 features. As Guo et al. (2020) states, the summed value at 80% of the rise time is used because it could be a valuable indication of the terrain's stiffness.

Temperature: the mean and variance of the terrain's temperature during every footstep are extracted for two features.

Classification Algorithms: We test four traditional machine learning algorithms and two deep learning algorithms to classify the extracted terrain features: k-nearest neighbor (KNN), support vector machine (SVM), random forest (RF), gradient boosting classifier (GB), artificial neural networks (ANN), and convolutional neural networks (CNN). Each classifier accepts the previously discussed extracted features as inputs for training. Techniques like K-Fold Cross Validation and Regression are implemented, as appropriate based on the algorithm, to find the best models.

4. EXPERIMENTAL SETUP

Finite Element Analyses (FEA) are performed on the servo linkages, the central part of the foot, and the servo bracket to validate the mechanical design of the foot. To simulate the forces acting on the foot, the bottom of the foot is fixed and a force is applied at the ankle joint equal to the weight of the robot. The servo bracket is fixed by the holes where the bracket is bolted to the central foot, and the force required to secure the servo at the servo's maximum torque is applied to the upper, inside surface of the servo bracket. For the linkages, hinge constraints are applied to the inner cylindrical surfaces where they rotate as the servo actuates, and additional constraints are applied to represent the bolt fixing the linkages at the linkage joint above the servo. The force acting on the linkage is applied to the inside face of the upper hole parallel to the linkage pointing towards the opposite hole.

Fig. 4 shows the result of the FEA on the linkage as an example. Darker blue areas represent areas under lower stress, while red areas represent those with high stress. Because FEA results show that the linkages have the most potential for failure, an endurance limit study, which is related to fatigue failure, is performed for the linkages using Marin Factors. The results of these analyses are in Section 5.

To gather training data for classification, we fabricate a prototype and install it on an Instron 5965 machine, which can be programmed to simulate a footstep. For each step, the machine lowers the foot onto the terrain until 300 N

Table 1. Hyperparameter Tuning

Algorithm	Hyperparameters	
KNN	# of Neighbors	10
	Leaf Size	30
	Weight Function	Uniform
	Distance Computation	Minkowski
SVM	Regularization Scaler	1.0
	Kernel	RBF
	Decision Function	One-VS-Rest
RF	# of Trees	100
	Maximum Depth	None
	Split Criteria	Gini Impurity
GB	# Loss Function	Logistic Loss
	Learning Rate	0.1
	# of Estimators	100
	Max Depth	3
ANN	Hidden Layers	2
	Nodes/Hidden Layer	50, 100
	Batch Size	32
	Activation Function	ReLU
	Optimizer	Adam
	Learning Rate	Adaptive
CNN	Convolution Layers	2
	Batch-Norm Layers	2
	Max Pooling Layers	2
	Fully Connected Layers	2
	Trainable Parameters	97,354
	Batch Size	32
	Optimizer	Adam
	Learning Rate	1×10^{-3}

of force is measured. The foot is held in this position for one second while the tarsal segments, which actuate based on inputs from the tactile sensors, lowered. The machine then returns the foot to a neutral position as the tarsal segments retract, where it rests for 0.25 seconds before the process repeats. To avoid settling, we regularly shuffle any loose materials that shift during data gathering to prevent erroneous readings. This cycle is repeated 1000 times for each terrain so that each terrain is represented by 1000 data points during classification.

Ten materials are used in the experiments to simulate real-world terrains: a metal plate (METAL), a wood board (WOOD), a foam block (FOAM), a yoga mat (MAT), a patch of artificial grass (GRASS), gravel (GRAVEL), pine straw (STRAW), concrete (CONCRETE), a patch of carpet (CARPET), and poppy seeds (POPPY). We select these materials because they roughly represent a range of terrains that a bipedal robot could encounter in indoor and outdoor deployments.

Once the training data is gathered, it is used to construct models for each of the algorithms discussed in Section 3.3. The 10,000 data points are divided and 80% is allocated to a training set and 20% is allocated to a test set. Each model is trained and tested on these same sets. Once these models are tuned and trained, they are saved and loaded onto the Raspberry Pi. To validate the feasibility of using our design in real-time applications, ten steps are timed and classified for each of the top three performing algorithms and the average classification speed is recorded. Table 1 outlines the set of hyperparameters that returns the highest test results for each model.
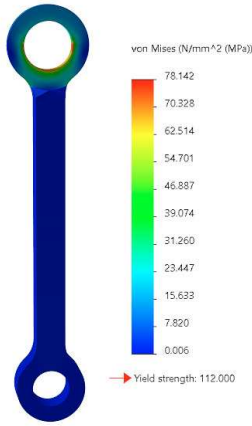
5. RESULTS AND DISCUSSION

Fig. 4. FEA results for the linkage.

FEA Analysis: FEA results for the central foot, linkages, and servo bracket can be seen in Table 2, which shows that the linkages have the most potential for failure with a factor of safety (FOS) of only 1.43. Despite the low FOS, the endurance limit study showed that the linkage design can withstand 5×10^8 footsteps before failing. At the speed that the Cassie robot walks, this comes out to constant walking for 24 hours per day for over six years. As such, the strength of these elements is validated to be reliable.

Learning Algorithm Evaluation: As shown in Table 4, RF has the highest test accuracy with 99.90%, while GB scores the second highest with 99.85% accuracy. Both deep learning-based models score nearly as well with CNN scoring 99.85% and ANN scoring 99.80%. This high performance could be due to the significant distinction in features generated per class. We complete a principal component analysis (PCA) to reduce the feature space to three dimensions for visualization purposes, which can be seen in Fig. 5. Each colored cluster represents the 100 features of each terrain projected onto three dimensions for visualization. The resulting feature space clearly shows each class in relatively dense clusters, excluding a small number of outliers. This visualization illustrates how the feature extraction discussed in Section 3.3 statistically groups the data from each terrain into unique clusters, making it simpler for the classification algorithms to identify the differences between each terrain. Based on the information obtained through the PCA, the high performance of both RF and GB also becomes clearer since tree-based models like RF and GB are known to outperform neural networks on tabular data, as seen in work by Grinsztajn et al. (2022) and Shwartz-Ziv and Armon (2022). SVM returned the lowest offline accuracy at 59.15%. However, based on the PCA visualization, this also may not be surprising considering that SVM works by linearly separating one class from the rest, and generally works better with two or three classes at the most.

Sensor Ablation Study: To understand how the classifiers learn, combinations of sensors are tested on RF and GB to see which sensors contribute the most to the classifiers. As Table 3 shows, each sensor had similar effects on RF and GB accuracy. However, the accelerometer and tactile sensor data influence the classification algorithms more than the data of the other sensors. While this shows the robustness of the classification algorithms

in being able to obtain sufficient information from only two sensors, it would be ideal if all of the sensors played a stronger role in the terrain identification to make the foot more resilient to failed or malfunctioning sensors. The reason for the disparity between sensors varies from sensor to sensor and could theoretically be caused by a number of reasons. More variation in the terrains, such as using hotter/colder and wet terrains, could increase the usefulness of the capacitive and temperature sensors. Moving the acoustic sensor away from the noisy servo could increase sound quality and make it more beneficial. And, as Section 3.3 outlines, 59 of the 100 features extracted from the raw data come from the tactile sensor. The classifiers might simply rely more on tactile information because they received more statistical data from that sensor. As such, expanding the feature extraction of the other sensors could increase their usefulness to the classifiers. It is expected that testing more diverse terrains (i.e., wet or cold terrains) will also increase the usefulness of the other sensors as the capacitive and temperature sensors would be more suitable than the accelerometer and tactile sensors to identifying these subtle differences. Beyond maximizing classification accuracy, using multiple sensors provides an additional benefit in terms of redundancy.
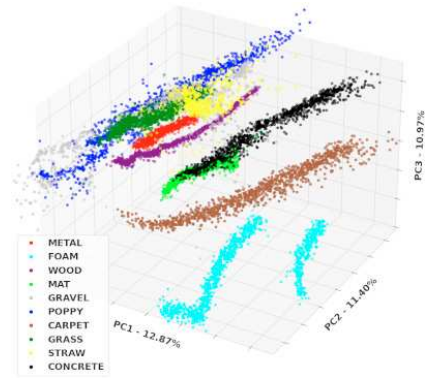

Fig. 5. Visualization of features projected to 3D space using PCA.

Table 3. Performance per Sensor

Sensor	RF Accuracy (%)	GB Accuracy (%)
all	99.90	99.85
temperature	49.45	54.73
accelerometer	96.17	96.07
microphone	60.10	61.09
capacitive	44.08	44.63
tactile	99.80	99.80
temp + mic + cap	86.22	85.67
temp + cap	69.40	71.19
temp + mic	82.84	83.08
mic + cap	68.26	68.71

Online Classification Speed: Models for GB, RF, and ANN are saved and loaded onto the Raspberry Pi to be timed. Although CNN outscores ANN by a slim margin, the accuracies are close enough that the ANN was used in place of the CNN because its architecture is more compatible with the microprocessor. Over ten steps, GB, RF, and ANN have average classification speeds of 2.73 ms, 32.76 ms, and 0.64 ms, respectively. RF has the longest inference time, possibly due to a large number of trees (see Table 1) within the model. ANN is the fastest by a large margin due to the relatively small size of the ANN model. Because classification is performed at the end of the footstep once the data for the step has been

Table 2. FEA Results

Component	Max Stress (MPa)	Yield Stress (MPa)	Factor of Safety
Central Foot	10.87	275.00	25.3
Linkages	78.14	112.00	1.43
Servo Bracket	57.61	112.00	1.94

Table 4. Test Accuracy Results (%)

	Metal	Foam	Wood	Mat	Gravel	Poppy	Carpet	Grass	Straw	Concrete	Average
RF	100.00	100.00	100.00	100.00	99.49	100.00	100.00	99.50	100.00	100.00	99.90
GB	100.00	100.00	100.00	100.00	99.49	100.00	99.53	99.50	100.00	100.00	99.85
CNN	100.00	100.00	100.00	100.00	98.98	100.00	100.00	99.50	100.00	100.00	99.85
ANN	100.00	100.00	100.00	100.00	100.00	99.52	100.00	99.50	99.04	100.00	99.80
k-NN	100.00	100.00	100.00	100.00	64.97	52.86	99.53	81.50	46.89	100.00	83.96
SVM	0.00	100.00	93.26	100.00	1.02	98.10	99.53	0.00	0.00	95.45	58.97

collected and processed, all of these times are fast enough that the data-gathering process of the following step would not be negatively impacted in a real-time application. This implies that the current system can realistically be used in an online terrain identification application.

6. CONCLUSIONS

Using multi-modal foot contact sensing consisting of an acoustic sensor, accelerometer, capacitive sensor, tactile sensor, and temperature sensor, the stabilizing Cassie foot was able to classify ten terrains with an average test accuracy of 99.90%, demonstrating great potential for real-world applications. To achieve the objective of using this foot in these scenarios, future work will be to implement a more powerful control scheme for tarsal segment actuation so the angle can be adaptively tuned to each terrain, and to assemble the foot onto the Cassie robot to obtain additional training data and complete real-time, online tests with the robot. Once the foot has successfully been integrated into the robot and tested, the classification algorithm will be used to update and optimize the robot's gait at runtime as it traverses a variety of terrains.

REFERENCES

Bednarek, M., Nowicki, M.R., and Walas, K. (2022). Haptr2: Improved haptic transformer for legged robots' terrain classification. *Robotics and Autonomous Systems*, 158, 104236.

Bosworth, W., Whitney, J., Kim, S., and Hogan, N. (2016). Robot locomotion on hard and soft ground: Measuring stability and ground properties in-situ. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 3582–3589. IEEE.

Catalano, M.G., Pollayil, M.J., Grioli, G., Valsecchi, G., Kolvenbach, H., Hutter, M., Bicchi, A., and Garabini, M. (2021). Adaptive feet for quadrupedal walkers. *IEEE Transactions on Robotics*, 38(1), 302–316.

Ding, L., Xu, P., Li, Z., Zhou, R., Gao, H., Deng, Z., and Liu, G. (2022). Pressing and rubbing: physics-informed features facilitate haptic terrain classification for legged robots. *IEEE Robotics and Automation Letters*, 7(3), 5990–5997.

Fondahl, K., Kuehn, D., Beinersdorf, F., Bernhard, F., Grimminger, F., Schilling, M., Stark, T., and Kirchner, F. (2012). An adaptive sensor foot for a bipedal and quadrupedal robot. In *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, 270–275.

Giguere, P. and Dudek, G. (2011). A simple tactile probe for surface identification by mobile robots. *IEEE Transactions on Robotics*, 27(3), 534–544.

Grinsztajn, L., Oyallon, E., and Varoquaux, G. (2022). Why do tree-based models still outperform deep learning on typical tabular data? In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Guo, X., Blaise, B., Molnar, J., Coholich, J., Padte, S., Zhao, Y., and Hammond, F.L. (2020). Soft foot sensor design and terrain classification for dynamic legged locomotion. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, 550–557.

Halatci, I., Brooks, C.A., and Iagnemma, K. (2008). A study of visual and tactile terrain classification and classifier fusion for planetary exploration rovers. *Robotica*, 26(6), 767–779.

Khan, Y.N., Komma, P., and Zell, A. (2011). High resolution visual terrain classification for outdoor robots. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 1014–1021. IEEE.

Khan, Y.N., Masselli, A., and Zell, A. (2012). Visual terrain classification by flying robots. In *International Conference on Robotics and Automation*, 498–503. IEEE.

Kõiva, R., Schwank, T., Walck, G., Meier, M., Haschke, R., and Ritter, H. (2020). Barometer-based tactile skin for anthropomorphic robot hand. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 9821–9826. IEEE.

Kolvenbach, H., Bärtschi, C., Wellhausen, L., Grandia, R., and Hutter, M. (2019). Haptic inspection of planetary soils with legged robots. *IEEE Robotics and Automation Letters*, 4(2), 1626–1632.

Kurobe, A., Nakajima, Y., and Saito, H. (2021). Audio-visual self-supervised terrain type recognition for ground mobile platforms. *IEEE Access*, 9, 29970–29979.

Piazza, C., Della Santina, C., Gasparri, G.M., Catalano, M.G., Grioli, G., Garabini, M., and Bicchi, A. (2016). Toward an adaptive foot for natural walking. In *IEEE-RAS 16th International Conference on Humanoid Robots*, 1204–1210.

Shamsah, A., Gu, Z., Warnke, J., Hutchinson, S., and Zhao, Y. (2023). Integrated task and motion planning for safe legged navigation in partially observable environments. *arXiv preprint arXiv:2110.12097*.

Shwartz-Ziv, R. and Armon, A. (2022). Tabular data: Deep learning is not all you need. *Information Fusion*, 81, 84–90.

Sójka, D., Nowicki, M.R., and Skrzypczyński, P. (2023). Learning an efficient terrain representation for haptic localization of a legged robot. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 12170–12176. IEEE.

Tenzer, Y., Jentoft, L.P., and Howe, R.D. (2014). The feel of mems barometers: Inexpensive and easily customized tactile array sensors. *IEEE Robotics & Automation Magazine*, 21(3), 89–95.

Venâncio, M.M., Gonçalves, R.S., and da Costa Bianchi, R.A. (2020). Terrain identification for humanoid robots applying convolutional neural networks. *IEEE/ASME Transactions on Mechatronics*, 26(3), 1433–1444.

Vulpi, F., Milella, A., Marani, R., and Reina, G. (2021). Recurrent and convolutional neural networks for deep terrain classification by autonomous robots. *Journal of Terramechanics*, 96, 119–131.

Walas, K. (2015). Terrain classification and negotiation with a walking robot. *Journal of Intelligent & Robotic Systems*, 78(3), 401–423.

Walas, K., Kanoulas, D., and Kryczka, P. (2016). Terrain classification and locomotion parameters adaptation for humanoid robots using force/torque sensing. In *IEEE-RAS International Conference on Humanoid Robots*, 133–140. IEEE.

Wang, M., Wonsick, M., Long, X., and Padr, T. (2020). In-situ terrain classification and estimation for nasa's humanoid robot valkyrie. In *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 765–770. IEEE.

Wu, X.A., Huh, T.M., Mukherjee, R., and Cutkosky, M. (2016). Integrated ground reaction force sensing and terrain classification for small legged robots. *IEEE Robotics and Automation Letters*, 1(2), 1125–1132.

Zhang, H., Dai, X., Sun, F., and Yuan, J. (2016). Terrain classification in field environment based on random forest for the mobile robot. In *Chinese Control Conference*, 6074–6079. IEEE.