Multi-Terrain Classification for Legged Robots Using HistGradient Boosting Machine Learning Technique

Yash Vardhan[1], Jyotindra Narayan[2,3], and Achraf Jabeur Telmoudi[4]
[1]*Department of Mechanical Engineering, Thapar Institute of Engineering and Technology, Patiala, Punjab, India*
[2]*Department of Mechanical Engineering, Indian Institute of Technology Patna, Bihar, India*
[3]*Department of Computing, Imperial College London, UK*
[4]*LISIER Laboratory, Higher National Engineering School of Tunis, University of Tunis, Tunisia*
yashvardhan7749@gmail.com, n.jyotindra@gmail.com, achraf.telmoudi@ieee.org

***Abstract*—The increasing adoption of legged robots for applications such as search and rescue, environmental monitoring, and planetary exploration presents unique challenges in navigating diverse and complex terrains. Accurate terrain classification is crucial for adaptive locomotion and reliable performance, yet existing approaches often suffer from limited generalizability and accuracy due to the inherent variability of terrains and sensor noise. This study proposes a multi-terrain classification framework leveraging the HistGradient Boosting (HGB) machine learning technique to address these challenges. The system utilizes a force sensor and IMU data from a quadruped robot to extract meaningful features for robust classification. The HistGradient Boosting model achieved the highest accuracy of** 0.9931**, highlighting its superior classification performance compared to k-Nearest Neighbors (KNN), Support Vector Machine (SVM), and Random Forest (RF) with accuracy of** 0.9410, 0.9878**, and** 0.9729**, demonstrating its effectiveness in handling the complexities of multi-terrain environments.**

***Index Terms*—terrain classification, legged robots, HistGradient Boosting, sensor data, machine learning**

I. Introduction

Legged robots have transformed exploration and intervention capabilities under challenging environments where conventional wheeled or tracked vehicles face operational limitations [1]. These robots prove invaluable in various scenarios, from disaster response to extraterrestrial terrain navigation, offering superior mobility and adaptability [2], [3]. Their unique advantage lies in navigating challenging landscapes, including debris, steep slopes, and uneven surfaces, making them essential for operations requiring sophisticated locomotion and terrain management. The biological inspiration behind legged robot design enables exceptional agility and durability across diverse environments [2], [4]. Additionally, these platforms show promising applications in military reconnaissance, environmental surveillance, and humanitarian operations where terrain diversity presents navigational obstacles [5].

Various machine learning methodologies have proven effective for terrain classification in robotic systems, each with distinct strengths and weaknesses [6]. Researchers have extensively applied supervised learning approaches, including Support Vector Machines (SVMs), Random Forests, and Neural Networks for terrain identification [7]. While decision trees provide clear classification logic, ensemble techniques such as bagging and boosting enhance prediction accuracy by combining multiple simple models [8]. In the deep learning domain, Convolutional Neural Networks (CNNs) have demonstrated exceptional capability in extracting complex spatial patterns from sensor inputs [9]. Gaussian Processes offer probabilistic classifications with uncertainty estimates, enhancing robot decision-making capabilities [10]. Further, transfer learning approaches facilitate model generalization across different robot platforms and sensor configurations [11]. Combined classification strategies integrating multiple techniques have shown enhanced performance compared to single-model approaches [12]. This range of classification models addresses the complex challenge of accurately identifying terrain types from sensor data.

Prior terrain classification research has primarily explored different sensor types, including force sensors, Inertial Measurement Units (IMUs), visual systems, and tactile sensors [7]. Initial research relied on hand-crafted features and rule-based approaches, which showed limited transferability between robotic platforms. Some studies utilized force sensor measurements to identify terrain properties by examining ground reaction forces and pressure patterns [13]. Approaches based on IMU data analyze angular velocities and linear accelerations to determine terrain characteristics, showing moderate effectiveness in controlled settings [7]. Visual classification methods employing texture analysis with machine learning have been investigated, though they face challenges with computational demands and lighting sensitivity [14]. Traditional classification models encounter several key challenges: diminished accuracy in complex environments, poor cross-platform generalization, and substantial computational requirements [6]. Effective terrain interaction classification requires adaptive techniques capable of real-time sensor fusion and robust feature extraction [5]. Moreover, most existing methods lack thorough testing across varied terrain types, restricting their practical deployment in critical operations [8]. The natural variability in sensor readings and terrain properties creates significant obstacles in developing universally applicable terrain classification frameworks.

This study, for the first time, presents a gradient-boosting machine learning framework for multi-terrain classification, specifically engineered to overcome the above-mentioned limitations in robotic terrain recognition. Specifically, this study proposes a HistGradient Boosting framework that efficiently fuses force sensor and IMU data, achieving 99.31% accuracy—outperforming conventional models. The HGB model optimizes computational efficiency using histogram binning, ensuring robust, adaptive, and real-time multi-terrain classification for legged robots in diverse environments. The major contributions are as follows.

(i) To propose a HistGradient Boosting-based framework for multi-terrain classification in legged robots in diverse and dynamic environments.
(ii) To utilize force sensor and IMU data fusion to improve terrain classification accuracy.
(iii) To optimize computational efficiency by leveraging histogram binning and gradient boosting techniques.
(iv) To provide a comparative analysis against conventional methods (SVM, KNN, and RF models), highlighting the advantages of ensemble learning for terrain recognition.

The structure of the paper is as follows. Section II presents the dataset details and sample data visualization and Section III mentions the proposed HGB method and performance metrics to assess its effectiveness. Section IV presents and discusses the numerical results while comparing them with state-of-the-art works. Finally, the concluding remarks are in Section V.

II. Dataset Details

The QCAT-legged robot dataset offers a detailed collection of sensor readings from the quadruped robot DyRET, specifically designed to study terrain interactions and locomotion efficiency under various conditions [15], [16]. A schematic representation of this is shown in Fig. 1. The dataset comprises two key sensor streams: force sensors positioned on each of the robot's four feet and an IMU mounted on its body. The force sensors capture triaxial force data (x, y, and z components) for each leg, providing insights into the robot-ground interaction. Meanwhile, the IMU records angular velocities, linear accelerations, and quaternion-based orientations, allowing precise tracking of the robot's movement in 3D space. The dataset encompasses six different terrains (such as concrete, gravel, and sand) and six-speed levels, with ten trials conducted for each terrain-speed combination. Each trial consists of eight steps, resulting in 2,880 recorded steps derived from the combination of terrains, speeds, trials, and steps. This extensive dataset is a valuable resource for robotic perception, control strategies, and terrain classification research.

Each trial in the dataset consists of approximately 5,300 samples spanning 8 steps, making it highly useful for machine learning applications. For visualization, a single step is illustrated in Fig. 2 for the IMU Sensor and Fig. 3 for the Force Sensor. This work used Jupyter notebooks to preprocess data from CSV files for terrain classification and gait optimization

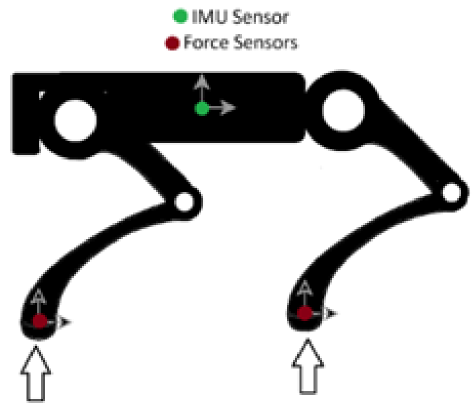

Fig. 1. Sensor placement and vector directions: IMU vector (green) shows body motion; force vectors (red) indicate ground reaction forces at the feet.

applications. The force sensor data captures the dynamic forces exerted by each leg across different surfaces and speeds, while the IMU sensor data provides insights into body dynamics under various locomotion scenarios. Studies utilizing this dataset have contributed to a deeper understanding of quadruped locomotion, including adaptive terrain navigation and gait stabilization [17]. With its well-structured yet diverse data, this dataset is crucial in advancing bio-inspired robotics and applying machine learning to robotic mobility.

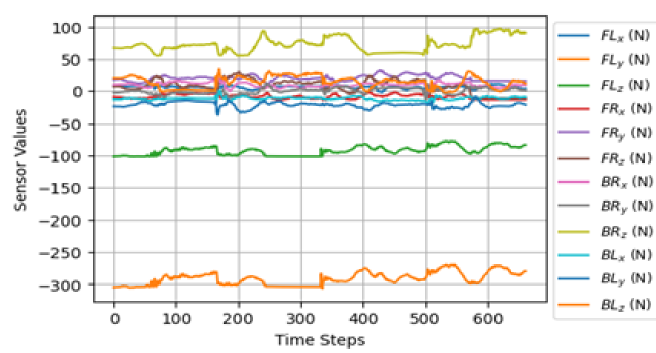

Fig. 2. Sample data visualization of the dataset of Force sensor readings for one step of the robot. The force sensors' values measure ground reaction forces (in N) on each leg (FL: Forward Left, FR: Forward Right, BR: Backward Right, BL: Backward Left) along x, y, z axes.

The computational resources used in this work include a multi-core CPU, AMD Ryzen 7 5700u, to train the HistGradient Boosting model efficiently. Since HGB is optimized for CPU computation, a GPU like an NVIDIA RTX 3060/3090 was unnecessary. The system had 16GB of RAM to handle the dataset and model training, along with an SSD (512GB) for fast data access. The implementation was performed in a Python environment using Jupyter Notebook or Google Colab, with machine learning libraries such as *scikit-learn* for model training, *numpy* and *pandas* for data processing, and *matplotlib* for visualization. CSV files were used to store and preprocess sensor data, while parallel processing was leveraged through *OpenMP* or *joblib* to speed up computations.

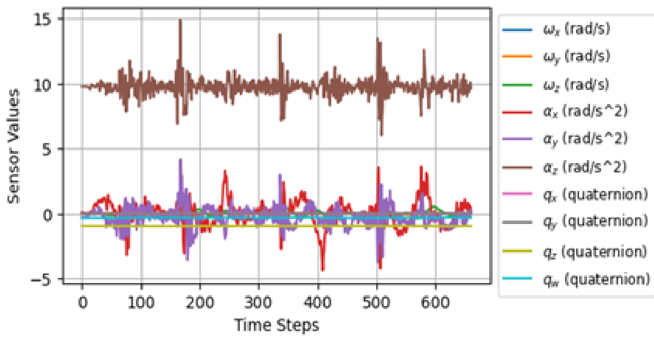

Fig. 3. Sample data visualization of the dataset of IMU sensor readings for one step of the robot. The IMU sensors' values capture angular velocity, ω (in rad/s), linear acceleration, α (in m/s²), and 3D orientation, q (in quaternion).

III. Methods and Metrics

This section presents a gradient-boosting model for multi-terrain classification for the quadruped robot and the performance metrics used to assess the model's effectiveness.

A. HistGradient Boosting Model

The proposed methodology for multi-terrain classification utilizes the HistGradient Boosting model, a powerful gradient-boosting technique that efficiently bins continuous feature values into histograms. This process allows the model to manage large-scale datasets with improved computational performance. HGB builds decision trees iteratively, minimizing residual errors from previous iterations to enhance predictive accuracy. Unlike traditional boosting methods, it does not require explicit memory storage of all sample gradients, making it computationally efficient and scalable.

HistGradient Boosting is a gradient boosting method that optimizes a loss function by sequentially fitting decision trees to the negative gradients of the loss function. It efficiently handles large datasets by binning continuous features into histograms, improving computational speed. The core idea of gradient boosting is to construct an additive model [18]:

$$F_m(x) = F_{m-1}(x) + \gamma h_m(x) \tag{1}$$

where $F_m(x)$ is the updated prediction function at iteration m, $h_m(x)$ is the weak learner (a decision tree) trained to minimize the residual loss, and γ is the learning rate controlling step size and scale the weak learner's contribution to prevent overfitting. Typical values of γ are between 0.01 and 0.2, ensuring gradual improvements.

At each iteration, the model minimizes a differentiable loss function $L(y, F(x))$, where y is the true target value and $F(x)$ is the predicted value. Instead of optimizing L directly, we compute the pseudo-residuals (negative gradients) for each sample:

$$r_i^{(m)} = -\left[\frac{\partial L\left(y_i, F\left(x_i\right)\right)}{\partial F\left(x_i\right)}\right]_{F(x)=F_{m-1}(x)} \tag{2}$$

This residual $r_i^{(m)}$ represents the direction in which the model should adjust its predictions to reduce the loss.

A weak learner $h_m(x)$, usually a shallow decision tree, is trained at step t to predict the pseudo-residual $r_i^{(m)}$, which effectively approximates the gradient of the loss function [19].

$$h_m(x) = \underset{h}{\operatorname{argmin}} \sum_{i=1}^{N} \left(r_i^{(m)} - h\left(x_i\right)\right)^2 \tag{3}$$

where N is the total number of training samples.

Unlike standard gradient boosting, HistGradient Boosting groups feature values into discrete bins, reducing memory consumption and improving efficiency. Instead of computing exact gradients for all samples, it approximates them using histogram binning:

$$g_k = \sum_{i \in B_k} g_i, \quad H_k = \sum_{i \in B_k} H_i \tag{4}$$

where g_k represents the summed gradient in k, H_k is the summed Hessian (second derivative of the loss function) in k, and B_k denotes all samples that fall into bin k.

HistGradient Boosting incorporates several regularization techniques. L1/L2 Regularization is applied to leaf values to prevent overfitting. Max Depth and Min Samples Split controls tree complexity to avoid overly deep trees. Early Stopping stops training when additional iterations do not significantly improve validation loss.

Finally, after M iterations, the final model prediction is given by:

$$\hat{y} = F_M(x) = F_0(x) + \sum_{m=1}^{M} \gamma h_m(x) \tag{5}$$

where $F_0(x)$ is the intial prediction.

This approach results in a model that efficiently approximates the optimal function mapping a to y_l, balancing accuracy and computational efficiency. A schematic architecture of the proposed HistGradient Boosting model is shown in Fig. 4.

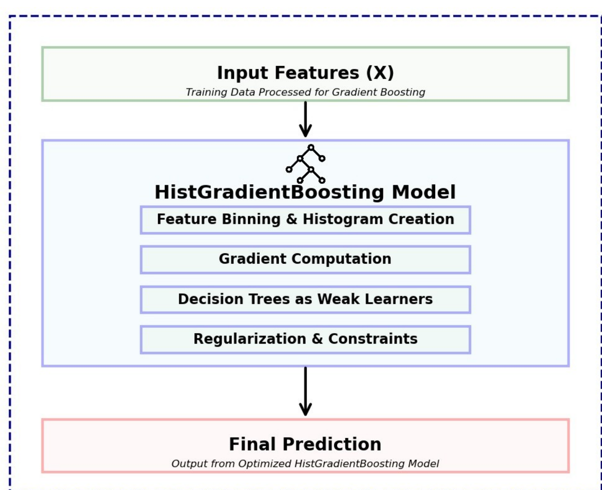

Fig. 4. HistGradient Boosting Model Architecture

B. Performance Metrics

Performance metrics are essential to assess the machine learning models in terrain classification. We evaluate using accuracy, precision, recall, F1-score, and the confusion matrix to capture different aspects of performance. While accuracy reflects the overall correctness, it can be misleading with imbalanced data. Hence, additional metrics ensure the classifier performs reliably across all terrain types. Precision shows how well the model avoids false positives, recall reflects its ability to minimize false negatives, and the F1-score balances both. These metrics are particularly useful for understanding how well the model distinguishes between terrains with overlapping features, such as gravel and sand.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{6}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{7}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{8}$$

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{9}$$

In the above expressions, TP, TN, FP, and FN represent True Positives, True Negatives, False Positives, and False Negatives, respectively. The confusion matrix provides a detailed breakdown of the model's predictions for each class. It highlights the TP, TN, FN, and FP for each terrain type, offering a visual and quantitative understanding of classification errors. By analyzing the confusion matrix, we can identify specific terrains where the model may struggle, such as misclassifying grass as dirt due to similar sensor patterns.

IV. Results and Discussions

Figure 5 revealed several important trends that guided the selection of an optimal configuration of the HistGradient Boosting model. When analyzing the effect of the learning rate ($learning\_rate$), it was observed that very low values (e.g., 0.01) resulted in slower convergence and lower test accuracy. As the learning rate increased to 0.1, there was a significant improvement in test accuracy, indicating a more effective learning process. Beyond this point, further increases in the learning rate led to only marginal improvements and introduced the risk of overfitting, as the training accuracy remained high while test accuracy plateaued or slightly declined. Similarly, the number of boosting iterations (max_iter) played a key role in refining model performance. Increasing the number of iterations from 50 to 200 led to steady improvements in both train and test accuracy. However, the gains in test accuracy started to diminish beyond 200 iterations, and further increases up to 300 showed limited additional benefits. This suggests that while more iterations allow the model to learn more complex patterns, they also contribute to longer training times without significantly enhancing generalization. The influence of the maximum tree depth (max_depth) was also evident in the tuning results. Increasing depth from 2 to 4 slightly improved training accuracy but led to a growing disparity between training and test accuracy. This trend indicates that deeper trees increased the model's capacity to memorize training data but compromised its generalization to unseen data. Based on the tuning results, the best performance was obtained with a learning rate of 0.1, a maximum depth of 2, and 200 boosting iterations. This configuration yielded a test accuracy of 0.9931, indicating high predictive performance.

Table I shows the classification performance of the HistGradient Boosting model across six different classes. Precision, recall, and F1-score values are high for all classes, with perfect scores (1.000) observed for classes 0, 2, and 5. Class 3 has a slightly lower recall (0.9551), which reduces its F1-score to 0.9770. The number of samples per class (support) is fairly balanced, ranging from 89 to 105, which contributes to a reliable evaluation of the model's performance across different categories. The overall accuracy of the model is 99.31%, indicating that most predictions were correct. Macro and weighted averages for precision, recall, and F1-score all exceed 99%, highlighting consistent performance regardless of class distribution. The model took approximately 77.82 seconds to train. These results indicate that the HistGradient Boosting model performs well in multi-class classification, maintaining consistent accuracy and class-wise balance.

TABLE I
RESULTS OF HISTGRADIENT BOOSTING MODEL

Class	Precision	Recall	F1-Score	Support
0	1.0000	1.0000	1.0000	105
1	0.9703	1.0000	0.9849	98
2	1.0000	1.0000	1.0000	96
3	1.0000	0.9551	0.9770	89
4	0.9890	1.0000	0.9945	90
5	1.0000	1.0000	1.0000	98
Accuracy	0.9931	-	-	576
Macro Avg	0.9932	0.9925	0.9927	576
Weighted Avg	0.9932	0.9931	0.9930	576
Training Time (sec)	77.8218	-	-	-

Fig. 6 shows a bar chart of F1 scores for six classes (Concrete, Grass, Gravel, Mulch, Dirt, and Sand), with support values (number of instances per class) annotated above each bar. The F1 scores are close to 1.0, indicating high classification performance across all classes. Slight variations exist, with Mulch and Grass having slightly lower F1 scores. Overall, the model demonstrates strong predictive accuracy across all classes. The classification results indicate that HGB outperforms conventional tree-based models while maintaining computational efficiency. The confusion matrix in Fig. 7 highlights high precision across all terrain classes, confirming its robustness for real-world applications.

The evaluation of multiple machine learning models for terrain classification revealed varying accuracy, precision, recall, and F1-score levels, as shown in Fig. 8. The KNN model achieved an accuracy of 0.9410 with a quick training time of 0.0568 seconds, making it a fast but slightly less accurate

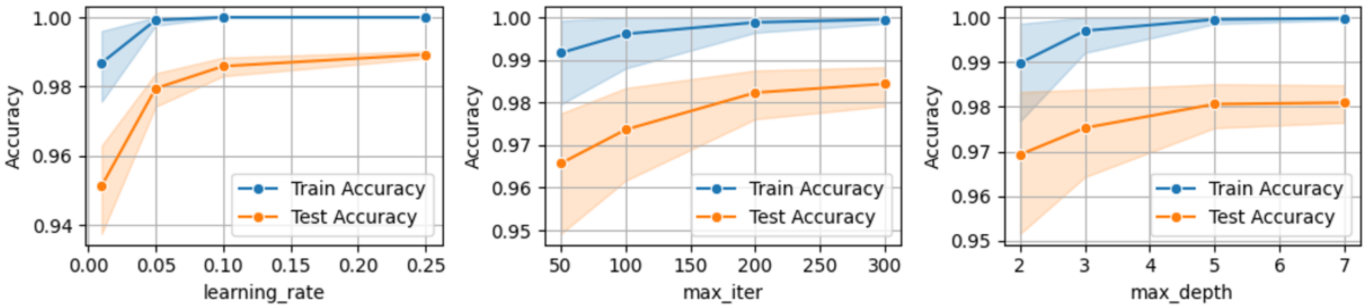

Fig. 5. Results of hyperparameter tuning of the HistGradient Boosting model: accuracy vs learning rate (left), accuracy vs maximum iterations (middle), and accuracy vs maximum depth (right), where shaded areas represent standard deviation or confidence intervals, across cross-validation folds.

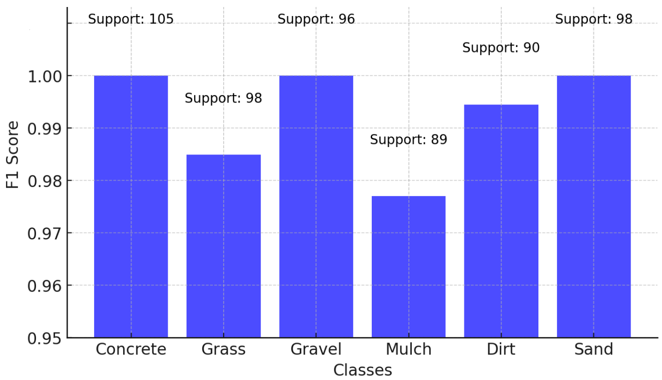

Fig. 6. F1 Score vs. Support for each class.

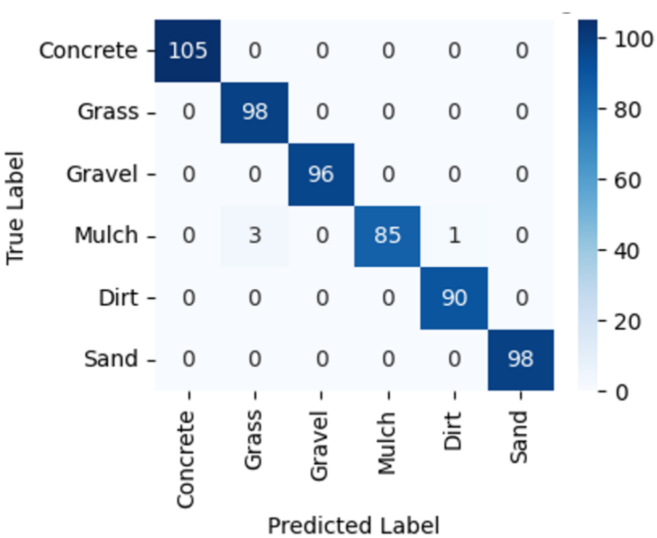

Fig. 7. Confusion matrix for HistGradient Boosting model

alternative. The SVM model significantly improved accuracy to 0.9878, with a well-balanced precision and recall, though at the cost of a longer training time of 24.7803 seconds. Random Forest performed slightly lower at 0.9792 accuracy, but its efficiency was notable with a training time of 21.2591 seconds. Gradient Boosting achieved an accuracy of 0.9861 but required an extensive training duration of 6845.5914 seconds, making it computationally expensive. HistGradient Boosting emerged as the best performer with an accuracy of 0.9931, striking a balance between accuracy and computational efficiency with a training time of 114.5423 seconds. These results indicate that while KNN offers the fastest training time and SVM delivers high accuracy with a trade-off in speed, HistGradient Boosting outperforms all models by achieving the highest accuracy while being significantly more efficient than traditional Gradient Boosting, making it the best overall choice for both performance and computational efficiency.

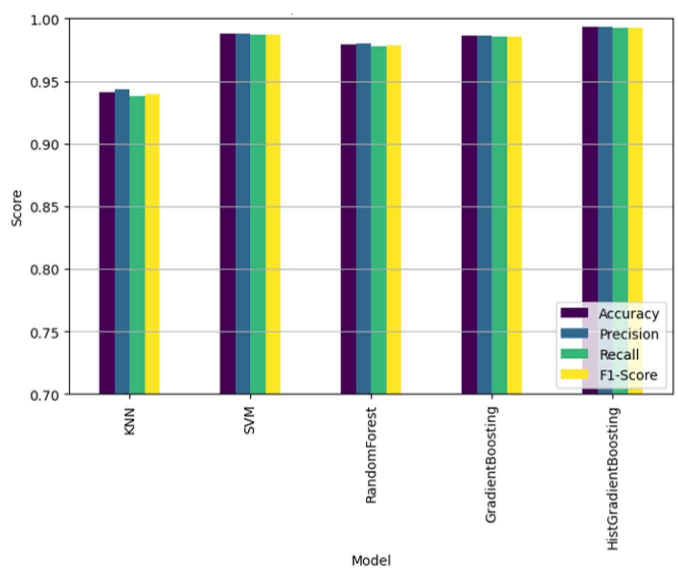

Fig. 8. Comparison of HistGradient Boosting model with different models

Our research findings demonstrated the superior performance of gradient-boosting approaches for multi-terrain classification in legged robotic systems. The HistGradient Boosting classifier achieved an exceptional accuracy of 0.9931, surpassing individual classification methods, including SVM (0.9878), KNN (0.9410), and Random Forest (0.9792). This section contrasts our results with prior studies that employed the QCAT dataset for terrain identification. In their research, Zhao et al. [20] implemented SVM algorithms for terrain classification in hexapod robots utilizing force-derived features from joint torques combined with IMU measurements. Their methodology facilitated adaptive movement across various surface conditions with a dynamic alternating tripod trotting gait. However, their SVM implementation reached only 96.7% accuracy, which falls below our HistGradient Boosting method's 99.31% accuracy, underscoring the benefits of ensemble techniques in enhancing classification results. Kumar et al. [21] investigated stacked LSTM-based loss regularization approaches for terrain classification using 89% of the QCAT dataset. While their methodology improved efficiency by eliminating preprocessing requirements for variable-length sequences, our gradient-boosting approach delivered superior classification accuracy. Ahmadi et al. [16] analyzed Gated Recurrent Neural Networks (GRNNs) and validated their models across both PUTany and QCAT datasets for terrain classification, determining that their semi-supervised approach effectively managed real-world terrain variations but demanded substantial training data for optimal performance. Compared to their RNN-based technique, our gradient boosting method achieved higher accuracy with reduced computational demands, making it more suitable for real-time robotic implementations.

Overall, this research contributes valuable insights to the terrain classification field by establishing that gradient boosting techniques can exceed the performance of both deep learning models and conventional clustering methods.

V. Conclusions

This research demonstrates the effectiveness of HistGradient Boosting in multi-terrain classification for legged robots. Its accuracy of 0.9931 offers a robust alternative to ensemble learning while significantly reducing training time. The study highlights that gradient-boosting techniques can achieve near-optimal classification without the high computational cost of stacking ensembles. The trade-off between accuracy and computational efficiency makes HGB an ideal candidate for deployment in real-time robotic applications. While the model has shown promising results, future research should focus on integrating additional sensor modalities, optimizing hyperparameter tuning strategies, and adapting the model for real-time inference in dynamic environments. Incorporating real-world terrain data and continuously updating the model using reinforcement learning approaches may further enhance its robustness and adaptability. By addressing these factors, HGB can contribute to advancing intelligent robotic mobility systems, ensuring reliable and adaptive terrain classification in diverse operational scenarios.

Acknowledgment

The authors acknowledge the computational resources offered by the Thapar Institute of Engineering and Technology, Patiala, which facilitated the training of this dataset.

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.

[2] Z. Yu, A. Duan, Z. Zhu, and W. Zhang, "Biomimetic soft-legged robotic locomotion, interactions and transitions in terrestrial, aquatic and multiple environments," *Sustainable Materials and Technologies*, p. e00930, 2024.

[3] B. Lindqvist, S. Karlsson, A. Koval, I. Tevetzidis, J. Haluška, C. Kanellakis, A.-a. Agha-mohammadi, and G. Nikolakopoulos, "Multimodality robotic systems: Integrated combined legged-aerial mobility for subterranean search-and-rescue," *Robotics and Autonomous Systems*, vol. 154, p. 104134, 2022.

[4] S. S. Kotha, N. Akter, S. H. Abhi, S. K. Das, M. R. Islam, M. F. Ali, M. H. Ahamed, M. M. Islam, S. K. Sarker, M. F. R. Badal *et al.*, "Next generation legged robot locomotion: A review on control techniques," *Heliyon*, 2024.

[5] M. Elnoor, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha, "Pronav: Proprioceptive traversability estimation for legged robot navigation in outdoor environments," *IEEE Robotics and Automation Letters*, 2024.

[6] M. Rybczak, N. Popowniak, and A. Lazarowska, "A survey of machine learning approaches for mobile robot control," *Robotics*, vol. 13, no. 1, p. 12, 2024.

[7] M. M. Kabir, J. R. Jim, and Z. Istenes, "Terrain detection and segmentation for autonomous vehicle navigation: A state-of-the-art systematic review," *Information Fusion*, vol. 113, p. 102644, 2025.

[8] M. Wang, L. Ye, and X. Sun, "Adaptive online terrain classification method for mobile robot based on vibration signals," *International Journal of Advanced Robotic Systems*, vol. 18, no. 6, p. 17298814211062035, 2021.

[9] R. Guo, W. Li, Y. He, T. Zeng, B. Li, G. Song, and J. Qiu, "Terrain slope parameter recognition for exoskeleton robot in urban multi-terrain environments," *Complex & Intelligent Systems*, vol. 10, no. 2, pp. 3107–3118, 2024.

[10] I. del Pino, A. Santamaria-Navarro, A. G. Zulueta, F. Torres, and J. Andrade-Cetto, "Probabilistic graph-based real-time ground segmentation for urban robotics," *IEEE Transactions on Intelligent Vehicles*, 2024.

[11] X. Yu, J. P. Queralta, and T. Westerlund, "Towards lifelong federated learning in autonomous mobile robots with continuous sim-to-real transfer," *Procedia Computer Science*, vol. 210, pp. 86–93, 2022.

[12] Z. Li, B. Li, Q. Liang, W. Liu, L. Hou, and X. Rong, "Research and realization of target following and autonomous obstacle avoidance algorithm of quadruped robot," in *2021 40th Chinese Control Conference (CCC)*. IEEE, 2021, pp. 3984–3989.

[13] C. Feng, K. Dong, and X. Ou, "A robot ground medium classification algorithm based on feature fusion and adaptive spatio-temporal cascade networks," *Neural Processing Letters*, vol. 56, no. 5, p. 235, 2024.

[14] J. Zhang, J. Guo, H. Chai, Q. Zhang, Y. Li, Z. Wang, and Q. Zhang, "A day/night leader-following method based on adaptive federated filter for quadruped robots," *Biomimetics*, vol. 8, no. 1, p. 20, 2023.

[15] R. Ahmadi, T. Nygaard, N. Kottege, D. Howard, and N. Hudson, "Qcat legged robot terrain classification dataset," 2020, accessed: 2025-03-19. [Online]. Available: https://doi.org/10.25919/5f88b9c730442

[16] A. Ahmadi, T. Nygaard, N. Kottege, D. Howard, and N. Hudson, "Semi-supervised gated recurrent neural networks for robotic terrain classification," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1848–1855, 2021.

[17] X. A. Wu, T. M. Huh, R. Mukherjee, and M. Cutkosky, "Integrated ground reaction force sensing and terrain classification for small legged robots," *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1125–1132, 2016.

[18] J. H. Friedman, "Greedy function approximation: a gradient boosting machine," *Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.

[19] E. Fouillen, C. Boyer, and M. Sangnier, "Proximal boosting: Aggregating weak learners to minimize non-differentiable losses," *Neurocomputing*, vol. 500, pp. 409–421, 2022.

[20] Y. Zhao, F. Gao, Q. Sun, and Y. Yin, "Terrain classification and adaptive locomotion for a hexapod robot qingzhui," *Frontiers of Mechanical Engineering*, vol. 16, no. 2, pp. 271–284, 2021.

[21] S. D. Kumar, S. Tripathi, K. Ujjwal, S. S. Jha, and S. De, "Robotic terrain classification with loss regularization," in *2023 3rd International Conference on Emerging Frontiers in Electrical and Electronic Technologies (ICEFEET)*. IEEE, 2023, pp. 1–6.