Proceedings of CONF-MLA 2026 Symposium: Intelligent Systems and Automation: AI Models, IoT, and Robotic Algorithms
DOI: 10.54254/2755-2721/2026.LD38755

A Safe Motion Control Method for Robots Based on Online Interpretable Decision Trees

Haoming Tang

Ulster College, Shaanxi University of Science and Technology, Xi'an, China
thm465478@163.com

Abstract. As one of the most interpretable machine-learning algorithms, decision trees have drawn increasing interest in intelligent robot control in recent years. As robots increasingly operate in dynamic and unstructured environments, the limitations of traditional black-box control strategies in interpretability, safety verification, and online adaptation have become increasingly prominent. Decision trees, with their inherently transparent reasoning structures and strong potential for formal verification, have emerged as ideal candidates for constructing trustworthy robotic controllers. This paper systematically reviews research on decision trees in robot control. Representative works are surveyed across five key dimensions: classical decision tree algorithms, online incremental decision trees, safety-constrained decision trees, interpretable reinforcement learning, and policy extraction. Furthermore, it analyzes the advantages and limitations of existing methods in robot motion control and outlines future research trends in this field.

Keywords: Decision trees, Robot control, Online learning, Explainable artificial intelligence (XAI), Safety verification

1. Introduction

Decision trees employ an "IF-THEN" rule structure, where each path from the root node to a leaf node corresponds to an explicit decision rule, collectively forming an interpretable set of rules. Consequently, they are inherently interpretable and offer significant advantages in decision-making scenarios that require human understanding and trust [1, 2]. However, as robotics transitions from structured factory environments to dynamic and open real-world settings, applications such as disaster rescue and domestic service place greater demands on control strategies. Although black-box methods—represented by deep reinforcement learning—have demonstrated strong capabilities in perception-to-action mapping, their behaviors remain uninterpretable, their safety boundaries are difficult to verify formally, and their sample efficiency in novel scenarios is exceptionally low, making it challenging to simultaneously achieve interpretability, safety verification, and rapid adaptation [3]. To this end, this paper systematically reviews robot control methods based on decision tree learning. We survey representative works across four dimensions: classical algorithms, online incremental learning, safety constraint injection, and policy extraction. Furthermore, we analyze the advantages, limitations, and applicable scenarios of these approaches, and outline future research directions.

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.LD38755

2. Classical decision tree algorithms and early applications in robotics

2.1. Overview of classical decision tree algorithms

The core concept of decision tree learning is to induce from data a tree-structured model capable of accurately classifying or regressing unseen samples. Classical decision tree algorithms primarily encompass three canonical paradigms: ID3, C4.5, and CART. The ID3 algorithm, introduced by Quinlan, uses information gain as the attribute selection criterion, recursively partitioning the sample space by selecting the attribute with the maximum information gain [1]. However, ID3 exhibits an inherent bias toward attributes with many values and cannot handle continuous attributes. As an improved iteration of ID3, the C4.5 algorithm substitutes information gain with the information gain ratio, while incorporating a discretization mechanism for continuous attributes alongside pruning strategies, thereby significantly enhancing generalization performance [4]. Unlike the information-theoretic ID3 and C4.5, the Classification and Regression Trees (CART) algorithm proposed by Breiman et al. employs the Gini index as its splitting criterion. CART is capable of handling both classification and regression tasks simultaneously, and its generated binary tree structure exhibits superior mathematical properties for formal analysis [2]. Specifically, the binary tree structure of CART ensures that every root-to-leaf path corresponds to a decision rule composed of multiple inequalities, a structured representation that is particularly advantageous for subsequent safety verification.

Although classical decision tree algorithms have achieved remarkable success in batch learning scenarios, they are fundamentally offline learning methods—requiring the prior collection of a complete dataset to construct the decision tree model in a single pass. While this limitation restricts their direct application in dynamic environments, it nevertheless has laid a solid theoretical foundation for the subsequent development of online decision trees.

2.2. Early explorations of decision trees in robotics

The application of decision trees to robot control can be traced back to the late 1990s. In 1998, Hamzei and Mulvaney proposed a self-organizing fuzzy decision tree method for online robot navigation control [5]. By combining fuzzy logic with decision tree structures, this approach enabled a robot to incrementally construct and refine its control strategy via online learning within unknown environments. As one of the representative early works applying decision trees to online robot control, this study primarily demonstrated that decision tree structures are suitable not only for offline classification tasks but also for adapting to environmental variations through incremental updates in real-time interactive scenarios. Although exploratory research during this period was limited in algorithmic maturity and experimental scale, it pioneered the extension of decision trees from static data analysis to dynamic robot control, paving the way for the subsequent development of online incremental decision trees.

3. Online incremental decision trees and real-time robot control

3.1. Hoeffding trees and data stream mining

In real-time robot control, data arrives continuously in a streaming fashion, posing immense computational and memory demands on traditional batch learning algorithms. Domingos and Hulten [6] introduced the Hoeffding Tree—also known as the Very Fast Decision Tree (VFDT)—in 2000,

DOI: 10.54254/2755-2721/2026.LD38755

extending decision tree learning to online streaming scenarios for the first time. Leveraging the Hoeffding bound, the algorithm determines the optimal splitting attribute once sufficient statistical confidence is attained, eliminating the need to wait for or store the entire dataset. Consequently, it processes high-speed data streams efficiently under resource-constrained settings while providing theoretical guarantees of asymptotic convergence to a tree learned via traditional batch processing.

Subsequently, Bifet and Gavaldà introduced drift detection mechanisms to handle concept drift [7]. Bifet et al. further improved the processing efficiency for high-dimensional data [8], while da Costa et al. reduced memory footprint without compromising accuracy [9]. Collectively, these advancements rendered online decision trees viable for deployment on resource-constrained robotic platforms.

3.2. Applications of online decision trees in robot navigation

The most direct application of online decision trees in robotics lies in real-time navigation and motion control. In unknown or dynamic environments, where a complete model cannot be acquired a priori, robots must rely on incoming sensor data streams to construct control strategies online. Hoeffding trees and their derivative algorithms can incrementally update policies, enabling robots to refine their behaviors while accumulating experience [5].

Notably, the state space in robot navigation is typically continuous and high-dimensional, whereas traditional decision trees require continuous attributes to be discretized, which can lead to information loss and degraded control precision. Consequently, practical implementations must integrate appropriate feature representation methods to map high-dimensional continuous state spaces into feature spaces that decision trees can process effectively—a crucial step that directly determines the upper bound on control performance.

Furthermore, the data distribution during online learning inherently drifts as the robot's behavior evolves. Ensuring the sustained efficacy of decision trees under such non-stationary conditions represents a major challenge in ongoing research.

4. Interpretable decision trees and policy distillation

4.1. Interpretability advantages of decision trees

Explainable Artificial Intelligence (XAI) has emerged as a central research theme in the broader field of AI. In their survey, Barredo Arrieta et al. systematically defined the conceptual framework of XAI, highlighting that interpretability is a critical attribute for constructing responsible and trustworthy AI systems [10]. In robotic control, interpretability extends beyond regulatory compliance, it directly impacts the safety and efficiency of human-robot collaboration—only when behaviors are fully understandable and predictable can human operators effectively supervise, intervene in, and place trust in autonomous systems.

The unique value of decision trees lies in the inherent transparency of their structure. While black-box models such as deep neural networks rely on massive parameter sets to construct non-linear mappings—rendering their decision logic largely incomprehensible to humans—each root-to-leaf path in a decision tree represents an explicit "condition-consequence" rule, collectively forming an interpretable rule set. This "IF-THEN" structure facilitates rule-by-rule inspection and comprehension, providing a direct foundation for safety assessment and fault diagnosis while enabling cross-validation against domain prior knowledge—a property of paramount practical

DOI: 10.54254/2755-2721/2026.LD38755

significance in safety-critical robotic applications. However, as a decision tree grows excessively large, its local interpretability remains intact, whereas its global interpretability diminishes.

4.2. Policy distillation from neural networks to decision trees

While decision trees possess clear interpretability advantages, their expressive capacity falls short of that of deep neural networks when handling high-dimensional perceptual inputs (e.g., images, point clouds) and complex continuous control tasks. To reconcile the strengths of both paradigms, researchers have proposed policy distillation—a framework that employs a pre-trained neural network as a "teacher" and extracts its knowledge to train a decision tree as a "student." Consequently, the resulting policy retains high performance close to that of the neural network while inheriting the transparent interpretability of decision trees.

Frosst and Hinton [11] first systematically explored distilling neural networks into soft decision trees, enhancing model interpretability by mimicking the soft output distribution of the teacher model. Bastani et al. [12] introduced VIPER, incorporating an adversarial retraining mechanism into policy distillation to actively sample states where the decision tree underperforms. By iteratively retraining on these critical states, VIPER effectively narrows the performance gap with the baseline neural network. Experimental results demonstrated that VIPER achieves near-optimal performance across multiple reinforcement learning benchmarks while providing a structured representation suitable for downstream safety verification. Extending this distillation paradigm to continuous action spaces, Paleja et al. [13] proposed interpretable continuous control trees, enabling decision trees to generate continuous control signals directly and significantly broadening their applicability to robotic control tasks.

5. Safety-constrained decision trees and formal verification

5.1. Formal embedding of safety constraints

In safety-critical robotic applications, control policies must provide reliable safety guarantees. Formal verification can offer stronger mathematical guarantees for system behavior than conventional testing, however, the non-linear and high-dimensional nature of deep neural networks makes their verification exceptionally difficult, and even undecidable in certain scenarios [14]. Owing to their structured rule representation, decision trees possess inherent advantages in formal verification.

Sun et al. proposed a formal verification method based on Satisfiability Modulo Convex (SMC) encoding to verify the safety of autonomous robots equipped with neural network controllers [3]. The dtControl framework introduced by Ashok et al. can systematically translate controllers with proven correctness (e.g., look-up table controllers) into compact decision tree representations, substantially enhancing interpretability while preserving correctness guarantees [9]. Furthermore, Zhu et al. presented a framework for verification, design, and adaptation aimed at ensuring end-to-end system-level safety, providing a viable pathway for safety guarantees in tasks such as robot navigation.

5.2. Safe decision-making under uncertainty

One of the primary challenges in robot motion control arises from uncertainty in environmental perception—factors such as sensor noise, occlusions, and dynamic obstacles inevitably induce state

DOI: 10.54254/2755-2721/2026.LD38755

estimation errors. Guaranteeing safe decision-making under uncertain conditions has thus emerged as a central research topic in safety-constrained decision trees. Existing studies have explored various approaches to integrate uncertainty modeling into decision tree frameworks. One strategy introduces probabilistic safety boundaries during training, enabling tree nodes to output a confidence level alongside action predictions and trigger a fallback to a conservative safety mode whenever the confidence falls below a predefined threshold. Another strategy leverages robust optimization to ensure that all possible states within the state estimation error bounds strictly satisfy safety constraints. The core rationale behind both paradigms is unified: exploiting the structural simplicity of decision trees to transform otherwise computationally intractable safety analyses into tractable, precisely solvable problems.

Nevertheless, uncertainty modeling in most current approaches remains relatively simplified. The joint handling of complex spatiotemporal correlations, multimodal uncertainty, and environmental dynamics remains an open problem.

6. Integration of differentiable decision trees and deep reinforcement learning

Researchers have sought to bridge the divide between decision trees and neural networks by proposing a class of "Differentiable Decision Trees" (DDTs). These methods aim to combine the end-to-end learning capabilities of neural networks with the intrinsic interpretability of decision trees.

Silva et al. [15] systematically investigated optimization methods for interpretable differentiable decision trees within the reinforcement learning (RL) paradigm. Traditional decision tree node splitting involves discrete choices—namely, selecting a specific attribute and a split threshold—which precludes the direct application of standard gradient descent algorithms. To address this, Silva et al. introduced a continuous relaxation technique for internal decision nodes, rendering the routing decisions at each node differentiable. Through this formulation, decision trees can be optimized end-to-end via policy gradient methods directly within an RL framework, achieving optimization efficiency comparable to that of neural networks while preserving the overall interpretable tree structure.

The significance of differentiable decision trees is that they enable closer integration between decision tree structures and deep reinforcement learning. On the one hand, differentiable decision trees can receive gradient signals online to perform incremental updates, thereby overcoming the adaptability limitations of conventional decision tree learning algorithms in non-stationary environments. On the other hand, once optimization is complete, the resulting decision tree retains its natural rule-based representation, which can be directly leveraged for downstream safety analysis and human-robot interaction. Consequently, a unified framework integrating "online learning, interpretability, and safety" becomes technically viable.

7. Conclusion

This paper presents a systematic review of intelligent robot control methods based on decision tree learning. Spanning from classical decision tree algorithms to online incremental decision trees, from policy distillation to safety-constraint injection, and from differentiable decision trees to their integration with deep reinforcement learning, research on decision trees in robotic control has evolved into a comprehensive methodological framework. Existing literature demonstrates that, owing to their inherently transparent reasoning structure, high potential for formal verification, and online incremental learning capability, decision trees possess unique advantages in constructing safe,

DOI: 10.54254/2755-2721/2026.LD38755

interpretable, and trustworthy robotic controllers. Looking ahead, as explainable artificial intelligence, formal verification, and deep reinforcement learning further converge, decision-tree-based robotic control methods are likely to be validated and deployed across an expanding range of real-world, safety-critical scenarios.

References

[1] Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning, 1*(1), 81-106.

[2] Breiman, L., Friedman, J., Olshen, R., & Stone, C. (1984). Classification and regression trees–crc press. Boca Raton, Florida, 685.

[3] Sun, X., Khedr, H., & Shoukry, Y. (2019, April). Formal verification of neural network controlled autonomous systems. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control* (pp. 147-156).

[4] Salzberg, S. L. (1994). C4. 5: Programs for machine learning by j. ross quinlan. morgan kaufmann publishers, inc., 1993.

[5] Hamzei, G. S., & Mulvaney, D. J. (1998, October). Self-organising fuzzy decision trees for robot navigation: An online learning approach. In *SMC'98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No. 98CH36218) (Vol. 3)* (pp. 2332-2337). IEEE.

[6] Domingos, P., & Hulten, G. (2000, August). Mining high-speed data streams. In *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 71-80).

[7] Bifet, A., & Gavalda, R. (2009, August). Adaptive learning from evolving data streams. In *International Symposium on Intelligent Data Analysis* (pp. 249-260). Springer Berlin Heidelberg.

[8] Bifet, A., Zhang, J., Fan, W., He, C., Zhang, J., Qian, J., ... & Pfahringer, B. (2017, August). Extremely fast decision tree mining for evolving data streams. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1733-1742).

[9] da Costa, V. G. T., de Leon Ferreira, A. C. P., & Junior, S. B. (2018). Strict very fast decision tree: A memory conservative algorithm for data stream mining. *Pattern Recognition Letters, 116*, 22-28.

[10] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82-115.

[11] Frosst, N., & Hinton, G. (2017). Distilling a neural network into a soft decision tree. arXiv preprint arXiv:1711.09784.

[12] Bastani, O., Pu, Y., & Solar-Lezama, A. (2018). Verifiable reinforcement learning via policy extraction. *Advances in Neural Information Processing Systems, 31*.

[13] Paleja, R., Niu, Y., Silva, A., Ritchie, C., Choi, S., & Gombolay, M. (2022). Learning interpretable, high-performing policies for continuous control problems. arXiv preprint arXiv:2202.02352, 10.

[14] Ashok, P., Jackermeier, M., Jagtap, P., Křetínský, J., Weininger, M., & Zamani, M. (2020, April). dtControl: Decision tree learning algorithms for controller representation. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control* (pp. 1-7).

[15] Silva, A., Killian, T., Jimenez, I., Son, S. H., & Gombolay, M. (2020, June). Optimization methods for interpretable differentiable decision trees applied to reinforcement learning. In *International Conference on Artificial Intelligence and Statistics* (pp. 1855-1865). PMLR.