Learning Interpretable Robot Policies from Demonstration

av

David Cáceres Domínguez

Akademisk avhandling

Avhandling för teknologie doktorsexamen i Computer Science,
som kommer att försvaras offentligt
fredag den 05 december 2025 kl. 09.00,
Hörsal T, Teknikhuset, Örebro universitet

Opponent: Professor Odest Chadwicke Jenkins
University of Michigan
United States of America

Örebro universitet
Institutionen för naturvetenskap och teknik
701 82 ÖREBRO

Abstract

The vision of robots assisting with our daily life tasks hinges on a fundamental challenge: most robot behaviors require programming by experts, creating a barrier for the many. Learning from Demonstration (LfD) is a family of methods with the potential to democratize robot programming by enabling laypeople to teach robots new skills simply by showing them. However, most LfD methods operate as *black-box* systems, making it difficult for humans to interpret, adapt, or reuse — key factors for human-robot collaboration and understanding.

In this thesis, we instead build upon Behavior Trees (BTs) — a transparent and manageable robot programming framework, whose reactive, modular design and potential for interoperability make them well-suited for constructing sophisticated robot behaviors. To bridge the gap between high-level planning and low-level control, engineers often embed entire behaviors within the BT's leaf nodes. While functional, this approach encapsulates complete skills, obscuring intermediate subgoals and undermining the transparency and modularity that BTs are meant to provide — ultimately limiting skill reuse and adaptability.

The structure of a BT plays a crucial role in effective behavior design. While several approaches have aimed to learn BT structures from demonstrations, they often rely on predefined action sets and state spaces. This necessity for expert-curated inputs constrains the robot's learning flexibility and reintroduces a degree of expert dependency. Compounding these challenges is a more foundational issue: the BT community lacks universally accepted definitions and rigorous evaluation methods for key properties such as *interpretability* and *modularity*. This absence leads to inconsistent claims and makes meaningful comparisons across different studies a significant challenge.

To address these gaps, this dissertation proposes a path toward more intuitive and effective robot learning, articulated through three key contributions:

First, this thesis formalizes core BT properties for robotics and introduces metrics for the systematic evaluation and comparison of learned policies. Building on this foundation, it investigates how different BT structures impact the interpretability of the control policy, identifying design patterns that best align with human intuition and understanding.

Second, this thesis presents a unified control framework that integrates BTs with a high-frequency Stack-of-Tasks (SoT) control strategy, enhancing transparency of BT policies by explicitly revealing the underlying subgoals. This approach allows BT nodes to function as hierarchical, high-frequency control objectives. Moreover, the resulting system achieves rapid reactivity in dynamic environments while supporting the coexistence of heterogeneous controllers and preserving a clean, modular decomposition of complex tasks.

Finally, this thesis introduces an end-to-end, label-free LfD pipeline that simultaneously learns the global BT structure and the underlying actions — modeled as Dynamic Movement Primitives — directly from raw demonstration data. By leveraging vision-language models to automatically extract and annotate state representations, this method eliminates the need for handcrafted action sets, predefined state spaces, and time-consuming manual labeling.

In summary, this thesis provides a comprehensive framework for learning interpretable, modular, and adaptable robot control policies from demonstration, bridging the gap between transparent policy representation and practical, high-frequency robot control, which marks a significant step toward making robot programming more accessible, robust, and understandable for both experts and non-experts.