Exploiting Dynamics Structures in Learning, Planning and Control for Effective Robot Autonomy

Thai Duong

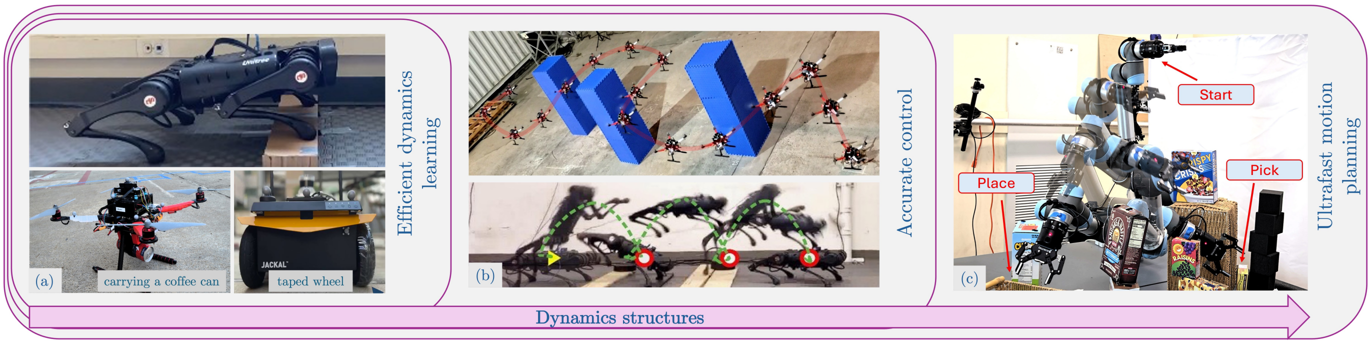

Fig. 1: My research exploits dynamics structures to achieve: (a) efficient dynamics learning of complex robot systems; (b) accurate trajectory tracking with the learned dynamics; and (c) ultrafast generation of dynamically feasible trajectories, e.g., for a pick and place task.

I. Motivation

The goal of my research is to develop effective, intelligent autonomous robot systems that work across a wide range of robots and scenarios, and self-improve through learning from past experiences in a human-centered environment. To build such robot systems for various dynamic and unstructured environments, e.g., homes, offices, warehouses, or hospitals, it is imperative to embrace both traditional planning and control techniques and machine learning advances. Robot planning and control rely on analytical models of the robot and the world to enable reasoning about what to do next (planning) and how to actually realize it (control). Meanwhile, learning enhances the robot's performance by obtaining hard-to-describe complex models and new skills from data. My goal is to develop *a principled framework for integrating learning with planning and control so that the robot can work efficiently, acquire and perfect its abilities via a continual self-improving loop of planning, control, and learning from past experiences*.

Towards this goal, my research seeks to unify machine learning, planning, and control techniques via prior structures, such as known physics knowledge and system properties. A key assumption in planning and control is the availability of an accurate robot dynamics model, governing the robot motion over time. While full-order complex motions are challenging to describe or learn from data, they are rich in dynamics structures, e.g., universal physics laws, that are known to be true and can be leveraged to assist learning. Such dynamics structures free learning from the burden of re-inferring known knowledge from data, and instead, redirect computational resources to train a lightweight residual dynamics model that can be updated easily over time from past experiences. I view prior dynamics structures as a unifying thread, stitching robot motions, machine learning, planning, and control together for an efficient autonomous system. Therefore, my work aims to:

- develop the theory, design, and realization of learning-based robot dynamics that preserve domain knowledge and prior structures by construction,
- develop efficient planning and control techniques with theoretical analysis based on the inherent structures,
- and enable resilient, long-term robot autonomy via a self-improving learning-planning-control-data loop.

Thus far, my research has focused on *preserving physics structures for efficient dynamic learning and accurate control* (Fig. 1(a)(b)), and *exploiting dynamics structures for ultrafast motion planning under dynamics constraints* (Fig. 1(c)), with extensive validation on various robot platforms, from mobile robots, legged robots, to manipulators (Fig. 1).

II. Past and Ongoing Research

Overview. My work leverages known physics structures of the robot dynamics and state spaces for efficient dynamics learning with long-term predictions, ultrafast motion planning and accurate control for safe and precise robot operations. Dynamics models from classical mechanics [1] are widely used in robotics, but may be over-simplified, leading to bias and modeling errors. Data-driven techniques [2–5] have emerged as a powerful approach to approximate robot dynamics, e.g., with neural networks, but typically require a large amount of data and training time. Hybrid approaches [6–9] have started combining physics knowledge with learning, but mainly focused on simple systems. Meanwhile, real and complex robot systems often describe their states on a Lie group manifold, e.g., robot/link poses on $SE(3)$, whose constraints are known a priori but pose a significant challenge to incorporate into a neural network. My work designs neural network architectures that capture both physics structures of the (continuous-time and discrete-time) robot dynamics, e.g., the law of energy conservation, and the manifold structures of the state space, e.g., a Lie group, for efficient training, generalization, accurate long-term predictions, and control designs. Furthermore, my work shows that other dynamics structures, such as differential flatness [10], can unlock the potential of motion planning under dynamics constraints by driving the planning times to the sub-millisecond range. My unified approach of learning, planning and control enables effective robot autonomy with: data-efficient dynamics learning and accurate trajectory tracking [11–14] under dynamics disturbances [15]; safety guarantees [16, 17]; accurate jumping and backflipping maneuvers [18, 19]; multi-agent tasks [20, 21], and ultrafast generation of dynamically feasible trajectories [22].

Learning and control of continuous-time dynamics. I proposed a *Hamiltonian-based neural ordinary differential equation (ODE) network on Lie groups* in [11–13, 15], whose architecture captures both the energy conservation law and the state space constraints via a Hamiltonian dynamics formulation on a Lie group [23]. The Hamilton's equations of motion and Lie group constraints are known physics structures, whose unknown quantities, such as kinetic energy, potential energy, energy dissipation, and input gains, are modeled by neural networks to capture hard-to-describe dynamics effects. The Hamiltonian structure allows us to design an energy-based controller [24] for trajectory tracking, successfully demonstrated on a quadrotor [12] (Fig. 1(a)(b)) and a non-holonomic ground robot [13] (Fig. 1(a)). While larger black-box models struggle to converge, my approach learns accurate dynamics from only a few minutes of state-control trajectories and provides accurate state predictions over *horizons hundreds of times longer* than those from black-box models. Our model can be fine-tuned quickly to handle dynamics changes, e.g., extra payload or hardware upgrades. The encoded structures enable *Lyapunov stability analysis* of our controller, which are typically unavailable with black-box dynamics models.

Learning and control of discrete-time dynamics. Discrete-time dynamics models, e.g., for widely used model predictive control (MPC) [25], are typically obtained by discretizing the continuous-time dynamics using numerical integration [26], without preserving physics laws and state space constraints, leading to high accumulated prediction errors. In [14], I developed a variational integrator (VI) network that preserves the energy conservation law and Lie group constraints in the discrete-time dynamics via the Lagrange-d'Alembert principle [27]. Unknown quantities, such as mass, potential energy, and input gains, are represented by neural networks and trained to capture intricate nonlinear dynamics effects from data. The preserved physics structures effectively regulate long-term prediction errors for accurate MPC control, e.g., on a quadrotor [14], and challenging jumping and backflipping maneuvers with legged robots [18, 19] (Fig. 1)(b)).

Leveraging dynamics structures for ultrafast kinodynamic planning. Planning a trajectory that respects the robot dynamics, i.e., *kinodynamic planning* [28], is crucial for safe and accurate robot operations. For high-degree-of-freedom robots, such as manipulators, sampling-based kinodynamic planners are common but their planning time is hindered by intractable boundary value problems (BVPs) [29, 30] during tree expansion. Recently, advances in parallelization [31, 32] have reduced planning times to less than a millisecond, but are limited to geometric planning. My research addresses these challenges by exploiting a dynamics structure, called *"differential flatness,"* and develops an extremely fast kinodynamic planning technique for a broad class of robot platforms, from mobile robots [33–35], manipulators [10], slider-pusher systems [36, 37], to mobile manipulators [38–40]. In my recent work [22], this structure transforms nonlinear dynamics into equivalent linear systems, and simplifies kinodynamic planning by obtaining a time-parameterized solution of the BVPs. Such closed-form solutions can be checked for collisions at multiple time samples in parallel, offering *sub-millisecond planning times* across benchmarks and real experiments (Fig. 1(c)). Via differential flatness, my work broadly transforms any existing geometric planners into ultrafast kinodynamic versions with dynamical feasibility guarantees.

III. Future Directions

My work thus far is the beginning of a long-term agenda towards effective robot autonomy by unifying learning, planning and control. My future directions will prioritize: time-critical robotic missions, contact-rich applications, and semantic structures of the environment for planning and control.

The time-parameterization of our trajectories allows us to specify and synthesize **deadline-driven robotic missions**, that explicitly reason about task ordering, temporal dependencies, and completion deadlines. I will develop a **time-critical task and motion planning** framework under such temporal constraints, formulated as "durative actions" in Planning Domain Definition Language [41], or enforced in Mission-time Linear Temporal Logics [42]. The resulting task plan can be quickly grounded to a trajectory by our kinodynamic planners.

In human-centric environments, it is important to consider **robot interaction with objects** in planning and control. Such interaction might create **additional constraints**, e.g., keeping a cup of water upright, that can be converted to the flat state space with simpler dynamics. It might involve **object pushing**, which is shown to be differentially flat [36, 37] and therefore, compatible with my kinodynamic planners. For control, I will integrate **learned contact dynamics**, e.g., by exploiting linear complementarity structures [43, 44], with parallelism for fast and accurate sampling-based MPC controllers [45, 46].

Furthermore, human-centric environments such as homes, offices, and hospitals are rich in **spatial and semantic structures**, e.g., in a scene graph map [47, 48], that can be used to **enhance planning and control**, e.g., via "common sense" reasoning capability of a foundation model [49, 50]. I have started exploring this direction by using a large language model with a scene graph map to guide an A* planner [51].

Finally, my future directions aim to develop an intelligent robot system for human-centered environments that humans can trust, thanks to embedded prior structures, while remaining capable of self-improving over time based on past experience.

References

[1] Lennart Ljung. System identification. *Wiley Encyclopedia of Electrical and Electronics Engineering*, 1999.

[2] Yunlong Song, Angel Romero, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Reaching the limit in autonomous racing: Optimal control versus reinforcement learning. *Science Robotics*, 8(82):eadg1462, 2023.

[3] Tim Salzmann, Elia Kaufmann, Jon Arrizabalaga, Marco Pavone, Davide Scaramuzza, and Markus Ryll. Real-time neural mpc: Deep learning model predictive control for quadrotors and agile robotic platforms. *IEEE Robotics and Automation Letters*, 8(4):2397–2404, 2023.

[4] Davide Scaramuzza and Elia Kaufmann. Learning agile, vision-based drone flight: From simulation to reality. In *The International Symposium of Robotics Research*, pages 11–18. Springer, 2022.

[5] Antonio Loquercio, Elia Kaufmann, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Learning high-speed flight in the wild. *Science Robotics*, 6(59):eabg5810, 2021.

[6] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.

[7] Zhengdao Chen, Jianyu Zhang, Martin Arjovsky, and Léon Bottou. Symplectic recurrent neural networks. *International Conference on Learning Representations (ICLR)*, 2020.

[8] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Symplectic ODE-Net: learning Hamiltonian dynamics with control. In *International Conference on Learning Representations (ICLR)*, 2019.

[9] Michael Lutter and Jan Peters. Combining physics and deep learning to learn continuous-time dynamics models. *The International Journal of Robotics Research*, 42(3): 83–107, 2023.

[10] Richard M Murray, Muruhan Rathinam, and Willem Sluis. Differential flatness of mechanical control systems: A catalog of prototype systems. In *ASME International Mechanical Engineering Congress and Exposition*, 1995.

[11] Thai Duong and Nikolay Atanasov. Hamiltonian-based neural ODE networks on the SE(3) manifold For dynamics learning and control. In *Robotics: Science and Systems*, 2021.

[12] Thai Duong, Abdullah Altawaitan, Jason Stanley, and Nikolay Atanasov. Port-Hamiltonian neural ODE networks on Lie groups for robot dynamics learning and control. *IEEE Transactions on Robotics*, 40:3695–3715, 2024.

[13] Abdullah Altawaitan, Jason Stanley, Sambaran Ghosal, Thai Duong, and Nikolay Atanasov. Hamiltonian dynamics learning from point cloud observations for nonholonomic mobile robot control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[14] Valentin Duruisseaux, Thai P Duong, Melvin Leok, and Nikolay Atanasov. Lie group forced variational integrator networks for learning and control of robot systems. In *Learning for Dynamics and Control Conference (L4DC)*, pages 731–744. PMLR, 2023.

[15] Thai Duong and Nikolay Atanasov. Adaptive control of SE(3) hamiltonian dynamics with learned disturbance features. *IEEE Control Systems Letters*, 6:2773–2778, 2022.

[16] Zhichao Li, Thai Duong, and Nikolay Atanasov. Robust and safe autonomous navigation for systems with learned SE(3) Hamiltonian dynamics. *IEEE Open Journal of Control Systems*, 1:164–179, 2022.

[17] Zhichao Li, Thai Duong, and Nikolay Atanasov. Safe autonomous navigation for systems with learned se (3) hamiltonian dynamics. In *Learning for Dynamics and Control Conference*, pages 981–993. PMLR, 2022.

[18] Chuong Nguyen, Abdullah Altawaitan, Thai Duong, Nikolay Atanasov, and Quan Nguyen. Variable-frequency model learning and predictive control for jumping maneuvers on legged robots. *IEEE Robotics and Automation Letters*, 10(2):1321–1328, 2025.

[19] Scott Beck, Chuong Nguyen, Thai Duong, Nikolay Atanasov, and Quan Nguyen. High accuracy aerial maneuvers on legged robots using variational integrator discretized trajectory optimization. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[20] Eduardo Sebastián, Thai Duong, Nikolay Atanasov, Eduardo Montijano, and Carlos Sagüés. LEMURS: Learning distributed multi-robot interactions. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 7713–7719, 2023.

[21] Eduardo Sebastian, Thai Duong, Nikolay Atanasov, Eduardo Montijano, and Carlos Sagues. Physics-informed multiagent reinforcement learning for distributed multirobot problems. *IEEE Transactions on Robotics*, 41: 4499–4517, 2025.

[22] Thai Duong, Clayton W. Ramsey, Zachary Kingston, Wil Thomason, and Lydia E. Kavraki. Ultrafast sampling-based kinodynamic planning via differential flatness. *arXiv:*, 2026.

[23] Taeyoung Lee, Melvin Leok, and N Harris McClamroch. *Global formulations of Lagrangian and Hamiltonian dynamics on manifolds*. Springer, 2017.

[24] Arjan Van Der Schaft and Dimitri Jeltsema. Port-Hamiltonian systems theory: An introductory overview. *Foundations and Trends in Systems and Control*, 1(2-3), 2014.

[25] F. Borrelli, A. Bemporad, and M. Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

[26] John Charles Butcher. *Numerical Methods for Ordinary Differential Equations*. John Wiley & Sons, 2016.

[27] J. E. Marsden and M. West. Discrete mechanics and variational integrators. *Acta Numer.*, 10:357–514, 2001.

[28] Andreas Orthey, Constantinos Chamzas, and Lydia E. Kavraki. Sampling-based motion planning: A compar-

ative review. *Annual Review of Control, Robotics, and Autonomous Systems*, 7(1):285–310, July 2024.

[29] Yanbo Li, Zakary Littlefield, and Kostas E. Bekris. Asymptotically optimal sampling-based kinodynamic planning. *The International Journal of Robotics Research*, 35(5):528–564, 2016.

[30] Dustin J. Webb and Jur van den Berg. Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In *IEEE International Conference on Robotics and Automation*, 2013.

[31] Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fabio Ramos, et al. Curobo: Parallelized collision-free robot motion generation. In *IEEE International Conference on Robotics and Automation*, 2023.

[32] Wil Thomason, Zachary Kingston, and Lydia E. Kavraki. Motions in microseconds via vectorized sampling-based planning. In *IEEE International Conference on Robotics and Automation*, 2024.

[33] Daniel Mellinger and Vijay Kumar. Minimum snap trajectory generation and control for quadrotors. In *IEEE International Conference on Robotics and Automation*, 2011.

[34] Katherine Mao, Jake Welde, M. Ani Hsieh, and Vijay Kumar. Trajectory planning for the bidirectional quadrotor as a differentially flat hybrid system. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

[35] Sikang Liu, Nikolay Atanasov, Kartik Mohta, and Vijay Kumar. Search-based motion planning for quadrotors using linear quadratic minimum time control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems* , 2017.

[36] Jiaji Zhou, Yifan Hou, and Matthew T Mason. Pushing revisited: Differential flatness, trajectory planning, and stabilization. *The International Journal of Robotics Research*, 38(12-13):1477–1489, 2019.

[37] Tom Lefebvre, Sander De Witte, Thomas Neve, and Guillaume Crevecoeur. Differential flatness of slider–pusher systems for constrained time optimal collision free path planning. *Journal of Dynamic Systems, Measurement, and Control*, 145(6):061001, 2023.

[38] Chin Pei Tang, Patrick T Miller, Venkat N Krovi, Ji-Chul Ryu, and Sunil K Agrawal. Differential-flatness-based planning and control of a wheeled mobile manipulator—theory and experiment. *IEEE/ASME Transactions on Mechatronics*, 16(4):768–773, 2010.

[39] Jake Welde and Vijay Kumar. Coordinate-Free Dynamics and Differential Flatness of a Class of 6-DOF Aerial Manipulators. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

[40] Skylar X Wei, Peder Harderup, and Joel Burdick. Differential flatness and flatness inspired control of aerial manipulators based on Lagrangian reduction. *arXiv preprint arXiv:2111.01302*, 2021.

[41] Maria Fox and Derek Long. PDDL2.1: An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 20:61–124, 2003.

[42] Jianwen Li, Moshe Y Vardi, and Kristin Y Rozier. Satisfiability checking for mission-time LTL. In *International Conference on Computer Aided Verification*, 2019.

[43] Wanxin Jin, Alp Aydinoglu, Mathew Halm, and Michael Posa. Learning linear complementarity systems. In *Learning for Dynamics and Control Conference*, 2022.

[44] Alp Aydinoglu, Adam Wei, Wei-Cheng Huang, and Michael Posa. Consensus complementarity control for multicontact mpc. *IEEE Transactions on Robotics*, 40: 3879–3896, 2024.

[45] Grady Williams, Paul Drews, Brian Goldfain, James M. Rehg, and Evangelos A. Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.

[46] Bogdan Vlahov, Jason Gibson, Manan Gandhi, and Evangelos A. Theodorou. MPPI-Generic: A CUDA library for stochastic trajectory optimization, 2024. URL https://arxiv.org/abs/2409.07563.

[47] Lukas Schmid, Marcus Abate, Yun Chang, and Luca Carlone. Khronos: A unified approach for spatio-temporal metric-semantic SLAM in dynamic environments. In *Robotics: Science and Systems (RSS)*, Delft, Netherlands, July 2024. doi: 10.15607/RSS.2024.XX.081.

[48] Dominic Maggio, Yun Chang, Nathan Hughes, Matthew Trang, Dan Griffith, Carlyn Dougherty, Eric Cristofalo, Lukas Schmid, and Luca Carlone. Clio: Real-time task-driven open-set 3D scene graphs. *IEEE Robotics and Automation Letters*, 9(10):8921–8928, 2024. doi: 10.1109/LRA.2024.3451395.

[49] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of large language models. *ACM Transactions on Intelligent Systems and Technology*, 2023.

[50] Mohamed Aghzal, Erion Plaku, Gregory J Stein, and Ziyu Yao. A survey on large language models for automated planning. *arXiv preprint arXiv:2502.12435*, 2025.

[51] Zhirui Dai, Arash Asgharivaskasi, Thai Duong, Shusen Lin, Maria-Elizabeth Tzes, George Pappas, and Nikolay Atanasov. Optimal scene graph planning with large language model guidance. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14062–14069, 2024.