

Hamiltonian Dynamics Learning from Point Cloud Observations for Nonholonomic Mobile Robot Control

Abdullah Altawaitan Jason Stanley Sambaran Ghosal Thai Duong Nikolay Atanasov

Abstract—Reliable autonomous navigation requires adapting the control policy of a mobile robot in response to dynamics changes in different operational conditions. Hand-designed dynamics models may struggle to capture model variations due to a limited set of parameters. Data-driven dynamics learning approaches offer higher model capacity and better generalization but require large amounts of state-labeled data. This paper develops an approach for learning robot dynamics directly from point-cloud observations, removing the need and associated errors of state estimation, while embedding Hamiltonian structure in the dynamics model to improve data efficiency. We design an observation-space loss that relates motion prediction from the dynamics model with motion prediction from point-cloud registration to train a Hamiltonian neural ordinary differential equation. The learned Hamiltonian model enables the design of an energy-shaping model-based tracking controller for rigid-body robots. We demonstrate dynamics learning and tracking control on a real nonholonomic wheeled robot.

I. INTRODUCTION

Autonomous mobile robot navigation is playing an increasingly important role in industrial manufacturing, agriculture, transportation, and warehouse automation. Reliable robot use over long duration requires robustness to changing conditions. Autonomous navigation requires adapting robot’s control strategy due to its own dynamics which can be done using model-free [1] or model-based [2] control techniques.

In this paper, we focus on model-based control of a rigid-body robot for trajectory tracking. The challenge is to infer a robot dynamics model from sensor observations and design a control policy that can successfully execute reference trajectories from a motion planner. Traditionally, a dynamics model is derived from first principles and its parameters are optimized using system identification techniques [3]. Hand-designed models, however, have a limited set of parameters and may not be sufficiently expressive to capture variations in the system dynamics. Recently, machine learning techniques have excelled at learn robot dynamics from data [4]–[8] by optimizing a large number of parameters in an expressive model, e.g., neural networks. However, these techniques require large datasets of state trajectories to learn a good dynamics model. Physics-informed neural networks reduce the amount of required data by embedding prior knowledge, such as physics laws and kinematic constraints, in the model architecture, ensuring that these requirements are satisfied

We gratefully acknowledge support from NSF CCF-2112665 (TILOS).

The authors are with the Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093, USA, e-mails: {aaltawaitan, jtstanle, sghosal, tduong, natanasov}@ucsd.edu. A. Altawaitan is also affiliated with Kuwait University as a holder of a scholarship.



Fig. 1: Clearpath Jackal robot equipped with a LiDAR.

by design instead of being inferred from data. This leads to better generalization and data efficiency. For example, Lagrangian [9]–[13] and Hamiltonian formulations [14]–[22] have been used to design neural network models to approximate dynamics of mechanical systems, where the equations of motions are enforced in the neural network architecture. Another approach is to assume known nominal dynamics for the system and only learn residual dynamics, e.g., when system dynamics change due to model uncertainty or disturbances. The residual dynamics can be learned online explicitly [23]–[29], allowing adaptation of the controller to dynamics uncertainty, or implicitly by learning controller’s parameters [30], [31]. Many physics-informed dynamics learning approaches, however, assume perfect state estimation for supervised training, and hence, are prone to estimation errors or are reliant on motion capture systems, limiting their real-world applications. Instead, our approach learns the robot dynamics directly from sensor observations, while maintaining a Hamiltonian-based machine learning model for data efficiency, generalization, and control design.

Learning control policies from observations, such as images or point clouds, has been extensively studied in reinforcement learning. A common approach is to learn a latent state space, and train a control policy in the latent state space to maximize the reward or control performance. For example, Zhang et al. [32] learn a latent space for control purposes using bisimulation metrics, which measure the similarity between two latent states if they have similar cumulative future reward. Kamran et al. [33] use convolutional neural networks to map images to an embedded space and learn a tracking controller via an actor-critic approach. Tian et al. [34] learn a latent state representation and latent dynamics model from observations to achieve discrete-time linear quadratic Gaussian regulation. The authors establish a finite-sample guarantee for obtaining a near-optimal state representation and control policy. In contrast, our work uses a known state representation but learns non-linear control-affine dynamics with Hamiltonian structure. We use a cycle consistency loss

that relates motion prediction from the dynamics model with motion prediction from registration of point-cloud observations. Cycle-consistency has been exploited in computer vision as a self-supervised learning technique to learn a model using unlabeled data for image alignment [35], [36], image-to-image translation [37], structure from motion [38], [39], and space-time alignment in videos [40], [41]. Recently, cycle consistency between observations has been used to transfer a learned control policy across different domains [42] without considering robot dynamics.

In this work, we develop a cycle consistency approach for robot dynamics learning from sensor observations. Our approach predicts the robot motion using a Hamiltonian dynamics model and measures the motion error using point cloud registration. Our prior work [20], [43] developed a Hamiltonian neural ODE (HNODE) method for learning quadrotor dynamics and an energy-shaping control formulation for quadrotor trajectory tracking. In this paper, we apply the HNODE approach to learn nonholonomic dynamics of a ground robot and supervise the learning process directly from sensor observations. Given a control input, we predict the robot's next pose using the dynamics model and obtain an observation from the sensors. The observations at two consecutive steps should be consistent, e.g., features from the first observation, after transformation to the next frame, should appear in the second observation. Violations of this cycle-consistency principle create a loss function that we use to train the Hamiltonian dynamics model via a neural ODE framework [44]. We then develop an energy-shaping trajectory tracking controller for the learned dynamics and demonstrate our approach in simulated and real autonomous navigation experiments with nonholonomic wheeled robots.

II. PROBLEM STATEMENT

Consider a robot modeled as a single rigid body with pose represented by position $\mathbf{p} \in \mathbb{R}^3$ and rotation matrix $\mathbf{R} \in SO(3)$. Let $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3 \in \mathbb{R}^3$ be the rows of \mathbf{R} , and define the robot's generalized coordinates as $\mathbf{q} = [\mathbf{p}^\top \ \mathbf{r}_1^\top \ \mathbf{r}_2^\top \ \mathbf{r}_3^\top]^\top \in \mathbb{R}^{12}$. Denote the generalized velocity of the robot by $\zeta = [\mathbf{v}^\top \ \boldsymbol{\omega}^\top]^\top \in \mathbb{R}^6$, where $\mathbf{v} \in \mathbb{R}^3$ is the linear velocity and $\boldsymbol{\omega} \in \mathbb{R}^3$ is the angular velocity, expressed in the body frame. The evolution of the robot state $\mathbf{x} = [\mathbf{q}^\top \ \zeta^\top]^\top \in \mathbb{R}^{18}$ is governed by control input \mathbf{u} according to a continuous-time dynamics model:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)). \quad (1)$$

We assume that the robot is equipped with a distance sensor, such as depth camera or LiDAR, which provides point clouds $\mathcal{Y}_n = \{\mathbf{y}_{n,m} \in \mathbb{R}^3\}_{m=1}^M$ at discrete time steps t_n . Assuming that the robot dynamics model is unknown, our objective is to design a control policy using a dataset of control inputs and point-cloud observations such that the robot is able to track a desired reference trajectory.

Problem 1. Consider a rigid-body robot with unknown dynamics model in (1). Let $\mathcal{D} = \{t_n^{(i)}, \mathcal{Y}_n^{(i)}, \mathbf{x}_0^{(i)}, \mathbf{u}_n^{(i)}\}_{i=1, n=0}^{D, N}$ be a training dataset of D trajectories. Each trajectory i

consists of an initial state $\mathbf{x}_0^{(i)}$, control inputs $\mathbf{u}_n^{(i)}$ applied at times $t_n^{(i)}$ and held constant during $[t_n^{(i)}, t_{n+1}^{(i)})$, and point clouds $\mathcal{Y}_n^{(i)}$ received by the robot at $t_n^{(i)}$ for $n = 0, \dots, N$. Design a control policy $\mathbf{u} = \pi(\mathbf{x}, \mathbf{x}^*(t))$ using the data \mathcal{D} such that the robot is able to track a desired reference trajectory $\mathbf{x}^*(t) = [\mathbf{q}^*(t)^\top \ \zeta^*(t)^\top]^\top$ for $t \in [0, \infty)$.

III. LEARNING FROM OBSERVATIONS USING CYCLE CONSISTENCY

We develop a model-based approach to solve Problem 1. We first learn a model of the system from point-cloud observations and then design a control policy $\pi(\mathbf{x}, \mathbf{x}^*(t))$ based on the learned model. We use a deep neural network model \mathbf{f}_θ to approximate the unknown dynamics \mathbf{f} in (1). Training a continuous-time neural ODE model can be done via the adjoint method [44]. Instead of a dataset \mathcal{D} of observations and controls as in Problem 1, the usual formulation uses a dataset $\{t_n^{(i)}, \mathbf{x}_n^{(i)}, \mathbf{u}_n^{(i)}\}_{i=1, n=0}^{D, N}$ of states and controls. The parameters θ are determined by minimizing a loss function, measuring the distance between predicted and actual states:

$$\begin{aligned} \min_{\theta} \quad & \sum_{i=1}^D \sum_{n=1}^N \ell_s(\mathbf{x}_n^{(i)}, \tilde{\mathbf{x}}_n^{(i)}) + \ell_r(\theta) \\ \text{s.t.} \quad & \dot{\tilde{\mathbf{x}}}^{(i)}(t) = \mathbf{f}_\theta(\tilde{\mathbf{x}}^{(i)}(t), \mathbf{u}^{(i)}), \quad \tilde{\mathbf{x}}^{(i)}(t_0) = \mathbf{x}_0^{(i)} \\ & \tilde{\mathbf{x}}_n^{(i)} = \tilde{\mathbf{x}}^{(i)}(t_n), \quad \forall n = 1, \dots, N, \quad \forall i = 1, \dots, D, \end{aligned} \quad (2)$$

with the predicted states $\tilde{\mathbf{x}}_n^{(i)}$ obtained from an ODE solver:

$$\tilde{\mathbf{x}}_n^{(i)} = \text{ODESolver}(\mathbf{x}_0^{(i)}, \mathbf{f}_\theta, \mathbf{u}_0^{(i)}, \dots, \mathbf{u}_n^{(i)}, t_0^{(i)}, \dots, t_n^{(i)}).$$

The error ℓ_r is a regularization term for the network parameters θ , while the error ℓ_s can be defined as a sum of square distance functions in \mathbb{R}^3 for position, $SO(3)$ for orientation, and \mathbb{R}^6 for generalized velocity (e.g., see [20]). The gradient of the loss function with respect to θ is obtained by solving an adjoint ODE backwards in time [44]. In this work, we extend the neural ODE formulation by (i) choosing \mathbf{f}_θ to satisfy Hamilton's equations of motion, which guarantee energy conservation and kinematic constraints, and (ii) training the model parameters θ using point-cloud observations $\mathcal{Y}_n^{(i)}$ rather than system states $\mathbf{x}_n^{(i)}$.

A. Imposing Hamiltonian Dynamics

To ensure energy conservation, the structure of \mathbf{f}_θ can be designed to follow Hamiltonian dynamics [18], [20]:

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{bmatrix} = \mathcal{J}_\theta(\mathbf{q}, \mathbf{p}) \begin{bmatrix} \nabla_{\mathbf{q}} \mathcal{H}_\theta(\mathbf{q}, \mathbf{p}) \\ \nabla_{\mathbf{p}} \mathcal{H}_\theta(\mathbf{q}, \mathbf{p}) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{g}_\theta(\mathbf{q}) \end{bmatrix} \mathbf{u}, \quad (3)$$

where $\mathbf{p} = \mathbf{M}_\theta(\mathbf{q})\zeta$ is the system momentum, $\mathbf{M}_\theta(\mathbf{q})$ is a positive semidefinite generalized mass matrix, $\mathbf{g}_\theta(\mathbf{q})$ is an input gain matrix, $\mathcal{J}_\theta(\mathbf{q}, \mathbf{p})$ is a state interconnection matrix, and $\mathcal{H}_\theta(\mathbf{q}, \mathbf{p})$ is the Hamiltonian modeling the total energy of the system as a sum of kinetic energy $\mathcal{T}_\theta(\mathbf{q}, \mathbf{p}) = \frac{1}{2} \mathbf{p}^\top \mathbf{M}_\theta(\mathbf{q})^{-1} \mathbf{p}$ and potential energy $\mathcal{V}_\theta(\mathbf{q})$:

$$\mathcal{H}_\theta(\mathbf{q}, \mathbf{p}) = \mathcal{T}_\theta(\mathbf{q}, \mathbf{p}) + \mathcal{V}_\theta(\mathbf{q}). \quad (4)$$

The dynamics of the generalized velocity ζ are described by $\dot{\zeta} = \dot{\mathbf{M}}_\theta(\mathbf{q})\mathbf{p} + \mathbf{M}_\theta(\mathbf{q})\dot{\mathbf{p}}$. Eq. (3) and $\dot{\zeta}$ describe our

approximated dynamics function $\mathbf{f}_\theta(\mathbf{x}, \mathbf{u})$ in (2). To model energy dissipation (e.g., friction) and ensure the kinematic constraints of the orientation matrix \mathbf{R} , the matrix $\mathcal{J}_\theta(\mathbf{q}, \mathbf{p})$ in (3) needs to have the following structure [20]:

$$\mathcal{J}_\theta(\mathbf{q}, \mathbf{p}) = \begin{bmatrix} \mathbf{0} & \mathbf{q}^\times \\ -\mathbf{q}^\times^\top & \mathbf{p}^\times \end{bmatrix} - \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_\theta(\mathbf{q}, \mathbf{p}) \end{bmatrix}, \quad (5)$$

where the first matrix is skew-symmetric and the second matrix models energy dissipation with positive semidefinite $\mathbf{D}_\theta(\mathbf{q}, \mathbf{p})$. The operators, \mathbf{q}^\times and \mathbf{p}^\times , are defined as:

$$\mathbf{q}^\times = \begin{bmatrix} \mathbf{R}^\top & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{r}}_1^\top & \hat{\mathbf{r}}_2^\top & \hat{\mathbf{r}}_3^\top \end{bmatrix}^\top, \quad \mathbf{p}^\times = \begin{bmatrix} \mathbf{0} & \hat{\mathbf{p}}_v \\ \hat{\mathbf{p}}_v & \hat{\mathbf{p}}_\omega \end{bmatrix}, \quad (6)$$

where $(\hat{\cdot}) : \mathbb{R}^n \mapsto \mathfrak{so}(n)$ maps $\mathbf{x} \in \mathbb{R}^n$ to a skew-symmetric matrix $\hat{\mathbf{x}} \in \mathfrak{so}(n)$. This model is referred to as HNODE with $\mathbf{M}_\theta(\mathbf{q})$, $\mathcal{V}_\theta(\mathbf{q})$, $\mathbf{g}_\theta(\mathbf{q})$, $\mathbf{D}_\theta(\mathbf{q}, \mathbf{p})$ as neural networks. Instead, we introduce a nominal model in each of the learnable terms and only learn residual components as follows:

$$\begin{aligned} \mathbf{M}_\theta(\mathbf{q}) &= (\bar{\mathbf{L}}(\mathbf{q}) + \tilde{\mathbf{L}}_\theta(\mathbf{q}))(\bar{\mathbf{L}}(\mathbf{q}) + \tilde{\mathbf{L}}_\theta(\mathbf{q}))^\top, \\ \mathbf{D}_\theta(\mathbf{q}, \mathbf{p}) &= (\bar{\mathbf{\Lambda}}(\mathbf{q}) + \tilde{\mathbf{\Lambda}}_\theta(\mathbf{q}))(\bar{\mathbf{\Lambda}}(\mathbf{q}) + \tilde{\mathbf{\Lambda}}_\theta(\mathbf{q}))^\top, \\ \mathcal{V}_\theta(\mathbf{q}) &= \bar{\mathcal{V}}(\mathbf{q}) + \tilde{\mathcal{V}}_\theta(\mathbf{q}), \quad \mathbf{g}_\theta(\mathbf{q}) = \bar{\mathbf{g}}(\mathbf{q}) + \tilde{\mathbf{g}}_\theta(\mathbf{q}), \end{aligned} \quad (7)$$

where the terms with bar and tilde accents represent fixed nominal models and learnable residuals, respectively. The structure of $\mathbf{M}_\theta(\mathbf{q})$ and $\mathbf{D}_\theta(\mathbf{q}, \mathbf{p})$ is chosen to ensure that they remain positive semidefinite using lower-triangular matrices $\bar{\mathbf{L}}(\mathbf{q})$, $\tilde{\mathbf{L}}_\theta(\mathbf{q})$, $\bar{\mathbf{\Lambda}}(\mathbf{q})$, $\tilde{\mathbf{\Lambda}}_\theta(\mathbf{q})$.

B. Training from Observations

Next, we focus on reformulating dynamics learning problem in (2) to use point-cloud observations $\mathcal{Y}_n^{(i)}$ instead of states $\mathbf{x}_n^{(i)}$. While point-cloud observations are available directly from RGBD cameras and LiDAR sensors found on most mobile robots, the robot states need to be estimated via LiDAR-based odometry, visual-inertial odometry, or a motion capture system. State estimation introduces localization errors in the dynamics learning process or requires additional infrastructure in the case of motion capture.

To enable learning from point-cloud observations, we introduce an observation-space error function ℓ_o and observation predictions $\tilde{\mathbf{y}}_{n,m}^{(i)}$ in the dynamics learning problem:

$$\begin{aligned} \min_{\theta} \quad & \sum_{i=1}^D \sum_{n=1}^N \sum_{m=1}^M \ell_o(\mathbf{y}_{n,m}^{(i)}, \tilde{\mathbf{y}}_{n,m}^{(i)}) + \ell_r(\theta) \\ \text{s.t.} \quad & \tilde{\mathbf{y}}_{n+1,m}^{(i)} = \tilde{\mathbf{R}}_{n+1}^{(i)\top} (\tilde{\mathbf{R}}_{n+1}^{(i)} \tilde{\mathbf{y}}_{n,m}^{(i)} + \tilde{\mathbf{p}}_n^{(i)} - \tilde{\mathbf{p}}_{n+1}^{(i)}), \\ & \dot{\tilde{\mathbf{x}}}^{(i)}(t) = \mathbf{f}_\theta(\tilde{\mathbf{x}}^{(i)}(t), \mathbf{u}^{(i)}), \\ & \tilde{\mathbf{x}}^{(i)}(t_0) = \mathbf{x}_0^{(i)}, \quad \tilde{\mathbf{y}}_{0,m}^{(i)} = \mathbf{y}_{0,m}^{(i)}, \\ & \tilde{\mathbf{x}}_n^{(i)} = \tilde{\mathbf{x}}^{(i)}(t_n), \quad \forall n = 0, \dots, N, \quad \forall i = 1, \dots, D, \end{aligned} \quad (8)$$

where $\ell_o(\mathbf{y}_{n,m}^{(i)}, \tilde{\mathbf{y}}_{n,m}^{(i)})$ can be selected as the square Euclidean distance $\|\mathbf{y}_{n,m}^{(i)} - \tilde{\mathbf{y}}_{n,m}^{(i)}\|_2^2$. The correspondence $(\mathbf{y}_{n,m}^{(i)}, \mathbf{y}_{n+1,m}^{(i)})$ is obtained using point-cloud registration [45]. The motivation for using a loss function in the observation space, instead of recovering the robot pose first, is to enable more general robot observations (e.g.,

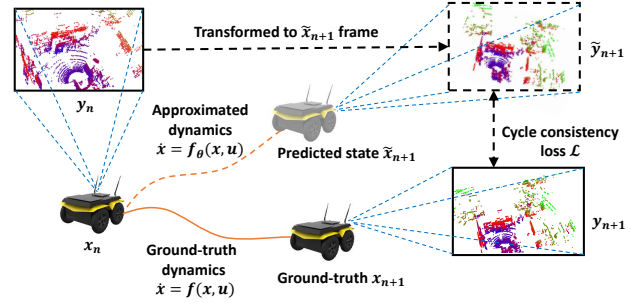


Fig. 2: Illustration of our observation-space loss function design.

image features) and states (e.g., velocity) in future work. The loss function can be modified to handle outliers, e.g., via a robust estimator, and even consider data association optimization. The key idea in this formulation is to relate motion predictions $\tilde{\mathbf{x}}_{n+1}^{(i)}$ from the learned dynamics with point-cloud prediction $\tilde{\mathbf{y}}_{n+1,m}^{(i)}$ in the first constraint in (8). As illustrated in Fig. 2, a point cloud $\mathcal{Y}_n^{(i)}$ observed at time t_n in the robot's body-frame is transformed to the world-frame using the predicted state $\tilde{\mathbf{x}}_n^{(i)}$. With the control input $\mathbf{u}_n^{(i)}$ at t_n , the model predicts the next state $\tilde{\mathbf{x}}_{n+1}^{(i)}$ at t_{n+1} , and transforms the 3D points back to the robot's body frame to obtain $\tilde{\mathcal{Y}}_{n+1}^{(i)}$. The discrepancy between these predictions and the actual observations is captured by ℓ_o . The term ℓ_r in (8) is used to promote sparsity in the learned dynamics model. For example, in underactuated systems the input gain \mathbf{g}_θ is often sparse because some inputs affect a subset of the states. Promoting sparsity through regularization ensures that the model remains simple and reflects the underactuation model structure accurately.

IV. CONTROL DESIGN FOR TRAJECTORY TRACKING

Given the Hamiltonian dynamics model learned from observations in Sec. III, a tracking control policy is designed using the interconnection and damping assignment passivity-based control (IDA-PBC) approach [46]–[48]. We present a general approach in Sec. IV-A and apply it to nonholonomic wheeled robots in Sec. IV-B.

A. IDA-PBC Control Approach

Let $\mathbf{x}^*(t) = [\mathbf{q}^*(t)^\top \boldsymbol{\zeta}^*(t)^\top]^\top$ be the desired trajectory that the system should track. For readability, we omit the time dependence of \mathbf{x}^* in the remainder of the paper. The desired momentum of the system in the body frame is calculated as $\mathbf{p}^* = \mathbf{M} \begin{bmatrix} \mathbf{R}^\top \mathbf{R}^* \mathbf{v}^* \\ \mathbf{R}^\top \mathbf{R}^* \boldsymbol{\omega}^* \end{bmatrix}$, leading to the desired $(\mathbf{q}^*, \mathbf{p}^*)$ for the learned Hamiltonian system. The Hamiltonian function $\mathcal{H}(\mathbf{q}, \mathbf{p})$ is generally not minimized at $(\mathbf{q}^*, \mathbf{p}^*)$. Therefore, the key idea of IDA-PBC is to inject energy to the system via the control input to achieve a desired Hamiltonian $\mathcal{H}_d(\mathbf{q}, \mathbf{p})$ that has a minimum at $(\mathbf{q}^*, \mathbf{p}^*)$. Let $(\mathbf{q}_e, \mathbf{p}_e)$ be the error in generalized coordinates and momentum:

$$\begin{aligned} \mathbf{R}_e &= \mathbf{R}^{*\top} \mathbf{R} = [\mathbf{r}_{e1} \quad \mathbf{r}_{e2} \quad \mathbf{r}_{e3}]^\top, \quad \mathbf{p}_e = \mathbf{p} - \mathbf{p}^*, \\ \mathbf{q}_e &= [\mathbf{p}_e^\top \quad \mathbf{r}_{e1}^\top \quad \mathbf{r}_{e2}^\top \quad \mathbf{r}_{e3}^\top]^\top, \quad \mathbf{p}_e = \mathbf{p} - \mathbf{p}^*. \end{aligned} \quad (9)$$

A common choice of $\mathcal{H}_d(\mathbf{q}, \mathbf{p})$ is:

$$\mathcal{H}_d(\mathbf{q}, \mathbf{p}) = \mathcal{T}_d(\mathbf{q}_e, \mathbf{p}_e) + \mathcal{V}_d(\mathbf{q}_e), \quad (10)$$

where $\mathcal{T}_d(\mathbf{q}_e, \mathbf{p}_e)$ and $\mathcal{V}_d(\mathbf{q}_e)$ are the desired kinetic and potential energy, respectively.

The IDA-PBC designs a control policy $\pi(\mathbf{x}, \mathbf{x}^*)$ such that the dynamics of the closed-loop system are also Hamiltonian with $\mathcal{H}_d(\mathbf{q}, \mathbf{p})$ as the total energy:

$$\begin{bmatrix} \dot{\mathbf{q}}_e \\ \dot{\mathbf{p}}_e \end{bmatrix} = \mathcal{J}_d(\mathbf{q}_e, \mathbf{p}_e) \begin{bmatrix} \frac{\partial \mathcal{H}_d}{\partial \mathbf{q}} \\ \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} \end{bmatrix}, \quad (11)$$

where $\mathcal{J}_d(\mathbf{q}_e, \mathbf{p}_e)$ is chosen in our design as follows [20]:

$$\mathcal{J}_d(\mathbf{q}_e, \mathbf{p}_e) = \begin{bmatrix} \mathbf{0} & \mathbf{J}_1 \\ -\mathbf{J}_1^\top & \mathbf{J}_2 \end{bmatrix} - \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{K}_d \end{bmatrix}, \quad (12)$$

$$\mathbf{J}_1 = \begin{bmatrix} \mathbf{R}^\top & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{r}}_{e1}^\top & \hat{\mathbf{r}}_{e2}^\top & \hat{\mathbf{r}}_{e3}^\top \end{bmatrix}^\top, \quad \mathbf{J}_2 = \mathbf{0}.$$

By matching the original dynamics in (3) and the closed-loop dynamics (11), we observe that the control input \mathbf{u} has to satisfy the following matching equations:

$$\mathbf{0} = \mathbf{J}_1 \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} - \mathbf{q}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} + \dot{\mathbf{q}} - \dot{\mathbf{q}}_e, \quad (13)$$

$$\mathbf{g}(\mathbf{q})\mathbf{u} = \mathbf{q}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{q}} - \mathbf{J}_1^\top \frac{\partial \mathcal{H}_d}{\partial \mathbf{q}_e} + \mathbf{J}_2 \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} - \mathbf{p}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} - \mathbf{K}_d \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} + \mathbf{D}_\theta(\mathbf{q}, \mathbf{p}) \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} + \dot{\mathbf{p}} - \dot{\mathbf{p}}_e. \quad (14)$$

Our choice of \mathbf{J}_1 and \mathbf{J}_2 in (12) satisfies the condition (13). We obtain the control policy $\mathbf{u} = \pi(\mathbf{x}, \mathbf{x}^*)$ from (14) as the sum $\mathbf{u} = \mathbf{u}_{ES} + \mathbf{u}_{DI}$ of an energy-shaping component \mathbf{u}_{ES} and a damping-injection component \mathbf{u}_{DI} :

$$\mathbf{u}_{ES} = \mathbf{g}^\dagger(\mathbf{q}) \left(\mathbf{q}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{q}} - \mathbf{J}_1^\top \frac{\partial \mathcal{H}_d}{\partial \mathbf{q}_e} + \mathbf{J}_2 \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} - \mathbf{p}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} + \mathbf{D}_\theta(\mathbf{q}, \mathbf{p}) \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} + \dot{\mathbf{p}} - \dot{\mathbf{p}}_e \right),$$

$$\mathbf{u}_{DI} = -\mathbf{g}^\dagger(\mathbf{q}) \mathbf{K}_d \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e}, \quad (15)$$

where $\mathbf{g}^\dagger(\mathbf{q}) = (\mathbf{g}^\top(\mathbf{q})\mathbf{g}(\mathbf{q}))^{-1} \mathbf{g}^\top(\mathbf{q})$ is the pseudo-inverse of $\mathbf{g}(\mathbf{q})$. Let $\mathbf{g}^\perp(\mathbf{q})$ be a maximal-rank left annihilator of $\mathbf{g}(\mathbf{q})$, i.e., $\mathbf{g}^\perp(\mathbf{q})\mathbf{g}(\mathbf{q}) = \mathbf{0}$. The control policy (15) exists when the following matching condition is satisfied:

$$\mathbf{g}^\perp(\mathbf{q}) \left(\mathbf{q}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{q}} - \mathbf{J}_1^\top \frac{\partial \mathcal{H}_d}{\partial \mathbf{q}_e} + \mathbf{J}_2 \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} - \mathbf{p}^\times \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} - \mathbf{K}_d \frac{\partial \mathcal{H}_d}{\partial \mathbf{p}_e} + \mathbf{D}(\mathbf{q}, \mathbf{p}) \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} + \dot{\mathbf{p}} - \dot{\mathbf{p}}_e \right) = 0. \quad (16)$$

B. Control Design for Nonholonomic Wheeled Robots

In this section, we consider an IDA-PBC control policy design for a nonholonomic wheeled robot. The wheels of a nonholonomic rigid-body only allow forces to be applied along the body's x -axis and torques around the z -axis. With two control inputs, we are able to control at most two degrees of freedom in the robot's configuration \mathbf{q} , making the system underactuated. The desired closed-loop dynamics (11) need

to be chosen carefully to ensure that there exists a control policy that can achieve them. Inspired by Lyapunov function designs in prior works [49], [50], where a nonholonomic wheeled robot is steered to approach a target pose at a specified angle, we design a desired Hamiltonian function that accounts for the underactuation. A natural choice of the desired kinetic and potential energies for a fully-actuated rigid-body robot is $\mathcal{T}_d(\mathbf{q}_e, \mathbf{p}_e) = \frac{1}{2} \mathbf{p}_e^\top \mathbf{M}(\mathbf{q}_e)^{-1} \mathbf{p}_e$ and $\mathcal{V}_d(\mathbf{q}_e) = \frac{k_p}{2} \mathbf{p}_e^\top \mathbf{p}_e + \frac{k_R}{2} \text{tr}(\mathbf{I} - \mathbf{R}_e)$, where $\mathcal{V}_d(\mathbf{q}_e)$ measures the quadratic Euclidean and chordal distance [51], respectively, from the current to the desired position and orientation with k_p and k_R as positive scalars. In the case of a nonholonomic wheeled robot, the orientation part of the desired potential energy needs to be shaped to guide the robot to approach the desired configuration \mathbf{q}^* at the correct orientation because the nonholonomic constraint prevents making sideways corrections instantaneously. We shape the desired potential energy as:

$$\mathcal{V}_d(\mathbf{q}_e) = \frac{k_p}{2} \mathbf{p}_e^\top \mathbf{p}_e + \mathcal{V}_{R_1}(\mathbf{q}_e) \mathcal{V}_{R_2}(\mathbf{q}_e) + \mathcal{V}_{R_3}(\mathbf{q}_e), \quad (17)$$

$$\text{where } \mathcal{V}_{R_1}(\mathbf{q}_e) = \frac{1}{2} \text{tr}(\mathbf{I} - \mathbf{R}_{\Delta_2}^\top \mathbf{R}(\mathbf{p}_e)^\top \mathbf{R}_{\Delta_1} \mathbf{R}^* \mathbf{R}_e),$$

$$\mathcal{V}_{R_2}(\mathbf{q}_e) = \frac{k_{R_1}}{2} \text{tr}(\mathbf{I} - \mathbf{R}^{*\top} \mathbf{R}_{\Delta_1}^\top \mathbf{R}(\mathbf{p}_e)), \quad (18)$$

$$\mathcal{V}_{R_3}(\mathbf{q}_e) = \frac{k_{R_2}}{2} \text{tr}(\mathbf{I} - \mathbf{R}(\mathbf{p}_e)^\top \mathbf{R}_{\Delta_1} \mathbf{R}^* \mathbf{R}_e).$$

The target-direction rotation matrix, $\mathbf{R}(\mathbf{p}_e)$, the rotations \mathbf{R}_{Δ_1} and \mathbf{R}_{Δ_2} are defined as:

$$\mathbf{R}(\mathbf{p}_e) = \begin{bmatrix} \frac{-\mathbf{p}_e}{\|\mathbf{p}_e\|} & \mathbf{U} \frac{-\mathbf{p}_e}{\|\mathbf{p}_e\|} & \mathbf{e}_3 \end{bmatrix}, \quad (19)$$

$$\mathbf{R}_{\Delta_1} = \mathbf{U} \exp\left(\frac{\hat{\boldsymbol{\theta}}_1 \pi}{\|\boldsymbol{\theta}_1\| 2}\right), \quad \mathbf{R}_{\Delta_2} = \exp\left(\frac{\hat{\boldsymbol{\theta}}_2 \pi}{\|\boldsymbol{\theta}_2\| 2}\right),$$

where $\mathbf{U} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$, $\boldsymbol{\theta}_1 = \log(\mathbf{U}^\top \mathbf{R}^{*\top} \mathbf{R}(\mathbf{p}_e))^\vee$,

$\boldsymbol{\theta}_2 = \log(\mathbf{R}^{*\top} \mathbf{R}_{\Delta_1}^\top \mathbf{R}(\mathbf{p}_e))^\vee$, and $(\cdot)^\vee : \mathfrak{so}(3) \mapsto \mathbb{R}^3$ is the inverse of the hat map. In (17), $\mathcal{V}_d(\mathbf{q}_e)$ is composed of three components. The first drives the robot position \mathbf{p} towards the desired position \mathbf{p}^* . The second approaches \mathbf{p}^* in a circular motion by orienting the robot perpendicular to the direction $-\mathbf{p}_e$ and aligning $\mathbf{R}(\mathbf{p}_e)$ with \mathbf{R}^* to specify the rotation direction. The third aligns \mathbf{R} with $\mathbf{R}(\mathbf{p}_e)$ for a direct straight-line path to \mathbf{p}^* . Once $\mathbf{R}(\mathbf{p}_e)$ is in alignment with \mathbf{R}^* , the first component becomes zero, leaving only the second component active.

Assuming a constant $\mathbf{M}_\theta(\mathbf{q}) \equiv \mathbf{M}_\theta$, Eq. (15) becomes:

$$\mathbf{u}_{ES} = \mathbf{g}(\mathbf{q})^\dagger \left(\mathbf{q}^\times \frac{\partial \mathcal{V}_\theta}{\partial \mathbf{q}} - \mathbf{p}^\times \mathbf{M}_\theta^{-1} \mathbf{p} + \mathbf{D}_\theta(\mathbf{q}, \mathbf{p}) \mathbf{M}_\theta^{-1} \mathbf{p} - \mathbf{e}(\mathbf{q}, \mathbf{q}^*) + \dot{\mathbf{p}}^* \right), \quad (20)$$

$$\mathbf{u}_{DI} = -\mathbf{g}(\mathbf{q})^\dagger \mathbf{K}_d \mathbf{M}_\theta^{-1} (\mathbf{p} - \mathbf{p}^*),$$

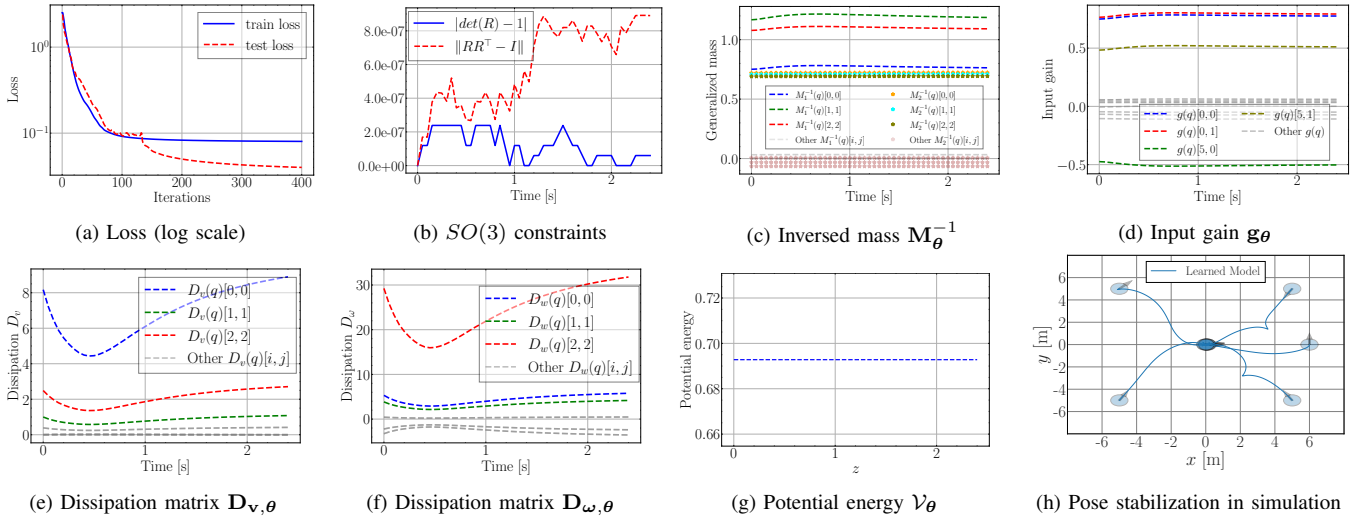


Fig. 3: Evaluation of our Hamiltonian Neural ODE network (a)-(g) along a trajectory and (h) pose stabilization in the Sapien simulator.

with $\mathbf{e}(\mathbf{q}, \mathbf{q}^*) = [\mathbf{e}_v(\mathbf{q}, \mathbf{q}^*)^\top \quad \mathbf{e}_\omega(\mathbf{q}, \mathbf{q}^*)^\top]^\top$ defined as:

$$\begin{aligned} \mathbf{e}_v(\mathbf{q}, \mathbf{q}^*) &= \mathbf{R}^\top \left(k_{\mathbf{p}} \mathbf{p}_e + \mathcal{V}_{R_2}(\mathbf{q}_e) \frac{\partial \mathcal{V}_{R_1}(\mathbf{q}_e)}{\partial \mathbf{p}_e} \right. \\ &\quad \left. + \mathcal{V}_{R_1}(\mathbf{q}_e) \frac{\partial \mathcal{V}_{R_2}(\mathbf{q}_e)}{\partial \mathbf{p}_e} + \frac{\partial \mathcal{V}_{R_3}(\mathbf{q}_e)}{\partial \mathbf{p}_e} \right), \quad (21) \\ \mathbf{e}_\omega(\mathbf{q}, \mathbf{q}^*) &= \frac{k_{R_2}}{2} \left(\mathbf{R}(\mathbf{p}_e)^\top \mathbf{R}_{\Delta_1} \mathbf{R} - \mathbf{R}^\top \mathbf{R}_{\Delta_1}^\top \mathbf{R}(\mathbf{p}_e) \right)^\vee \\ &\quad + \frac{\mathcal{V}_{R_2}(\mathbf{q}_e)}{2} \left(\mathbf{R}_{\Delta_2}^\top \mathbf{R}(\mathbf{p}_e)^\top \mathbf{R}_{\Delta_1} \mathbf{R} - \mathbf{R}^\top \mathbf{R}_{\Delta_1}^\top \mathbf{R}(\mathbf{p}_e) \mathbf{R}_{\Delta_2} \right)^\vee. \end{aligned}$$

Please refer to the extended version¹ of the paper [52] for the computation of the derivatives of the terms \mathcal{V}_{R_i} .

V. EVALUATION

In this section, we verify the effectiveness of our approach in both simulated and real experiments with a Jackal robot. We describe data collection for training in Sec. V-A, and then evaluate our training results and control performance in simulation (Sec. V-B) and in real experiments (Sec. V-C).

A. Data Collection

We collected a dataset $\mathcal{D} = \{t_n^{(i)}, \mathcal{Y}_n^{(i)}, \mathbf{x}_0^{(i)}, \mathbf{u}_n^{(i)}\}_{i=1, n=0}^{D, N}$ consisting of D N -sample trajectories with sampling interval of dt seconds by driving the robot manually using a joystick. The convergence of model parameter estimates requires a persistence of excitation condition to be satisfied [53]. Hence, it is important to collect a diverse range of linear and rotational motions. A short trajectory duration N is chosen to avoid numerical instability of the predicted states from an ODE solver. Each data point includes a point cloud $\mathcal{Y}_n^{(i)}$ from a 3D LiDAR sensor and the control inputs $\mathbf{u}_n^{(i)}$. From two consecutive LiDAR scans, we apply fast global registration [45] to pick M random correspondences to create two corresponding point clouds $\mathcal{Y}_n^{(i)}$ and $\mathcal{Y}_{n+1}^{(i)}$.

¹Paper website: <https://altwaitan.github.io/DLFO/>.

B. Sapien Simulation

We build a simulated differential drive robot in a physics-based simulation environment Sapien [54]. The control input is $\mathbf{u}_n^{(i)} = [\tau_L \quad \tau_R]^\top \in \mathbb{R}^2$, consisting of left τ_L and right τ_R wheel torques. A dataset \mathcal{D} was collected as described in Sec. V-A with $(N, D, M, dt) = (250, 32, 40, 0.05)$ by driving the robot using sampled wheel torques.

We train our dynamics model as described in Sec. III with the dataset \mathcal{D} for 400 iterations. The networks are pretrained to a nominal model $\mathbf{M}_0(\mathbf{q}) = \mathbf{I}$, $\mathbf{D}_0(\mathbf{q}) = 10^{-3} \mathbf{I}$, $\mathcal{V}_0(\mathbf{q}) = 0$, $\mathbf{g}_0(\mathbf{q}) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}^\top$ to avoid numerical instability by the ODE solver. The regularization loss $\ell_r(\boldsymbol{\theta})$ is set to $10^{-3} \|\mathbf{g}_\theta(\mathbf{q})\|_1$ to encourage sparsity in $\mathbf{g}_\theta(\mathbf{q})$. Fig. 3 shows training and testing performance for our HNODE design. The $SO(3)$ constraints are satisfied, by design, and the potential energy $\mathcal{V}_\theta(\mathbf{q})$ remains constant due to 2D plane motion. The learned mass $\mathbf{M}_\theta(\mathbf{q})$ converged to a diagonal matrix and the input gain $\mathbf{g}_\theta(\mathbf{q})$ is sparse with nearly equal values in its first row and opposite signs in the last, consistent with the nonholonomic constraint.

The controller proposed in Section IV-B was designed with the learned parameters and control gain values $(k_{\mathbf{p}}, k_{R_1}, k_{R_2}, \mathbf{K}_d) = (1.2, 7, 3, \text{diag}(1.2\mathbf{I}, \mathbf{I}))$. Fig. 3h shows that our controller successfully drives the robot to the desired pose from various initial poses with desired yaw achieved before converging to desired position.

C. Real World Deployment

We use a Clearpath Jackal robot to demonstrate our approach. The control input $\mathbf{u} = [\Delta\omega_1 \quad \Delta\omega_2]^\top \in \mathbb{R}^2$ represents the difference between the desired and actual wheel velocity $\Delta\omega_i = \omega_i^* - \omega_i$. The closed-loop dynamics of the wheels are modeled as a first order system, similar to the approach in [55], with $\dot{\omega} = \frac{1}{\alpha}(\omega^* - \omega)$, where α denotes the system's time constant. In doing so, we establish a means to control the wheel's angular acceleration, which directly correlates with torque, effectively satisfying Hamiltonian

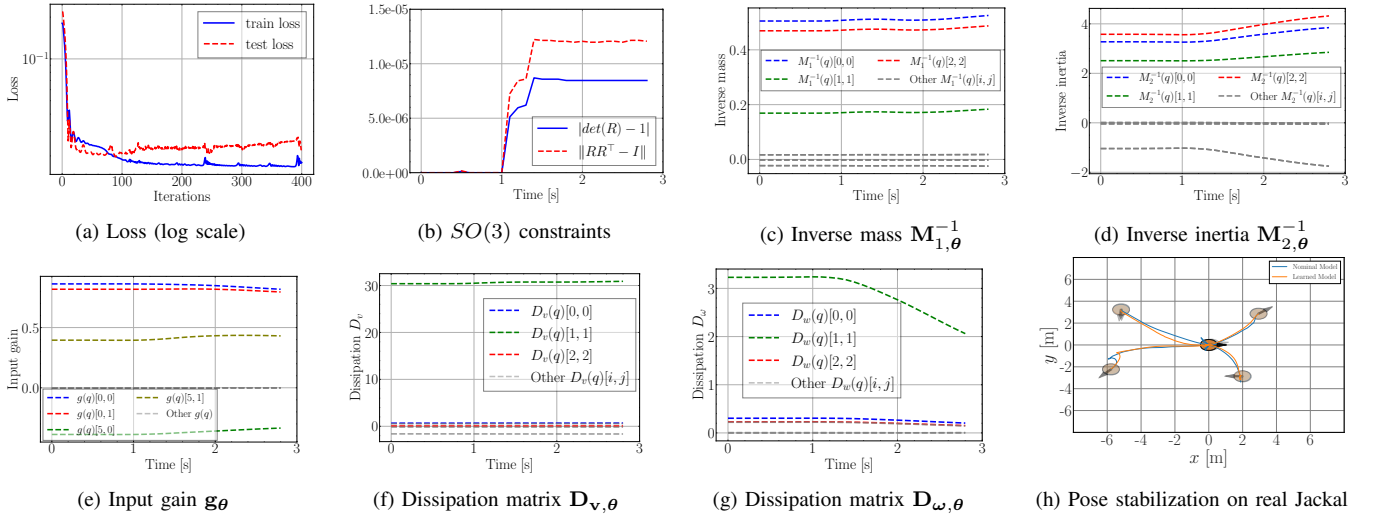


Fig. 4: Evaluation of our Hamiltonian neural ODE network (a)-(g) along a trajectory and (h) pose stabilization on a real Jackal.

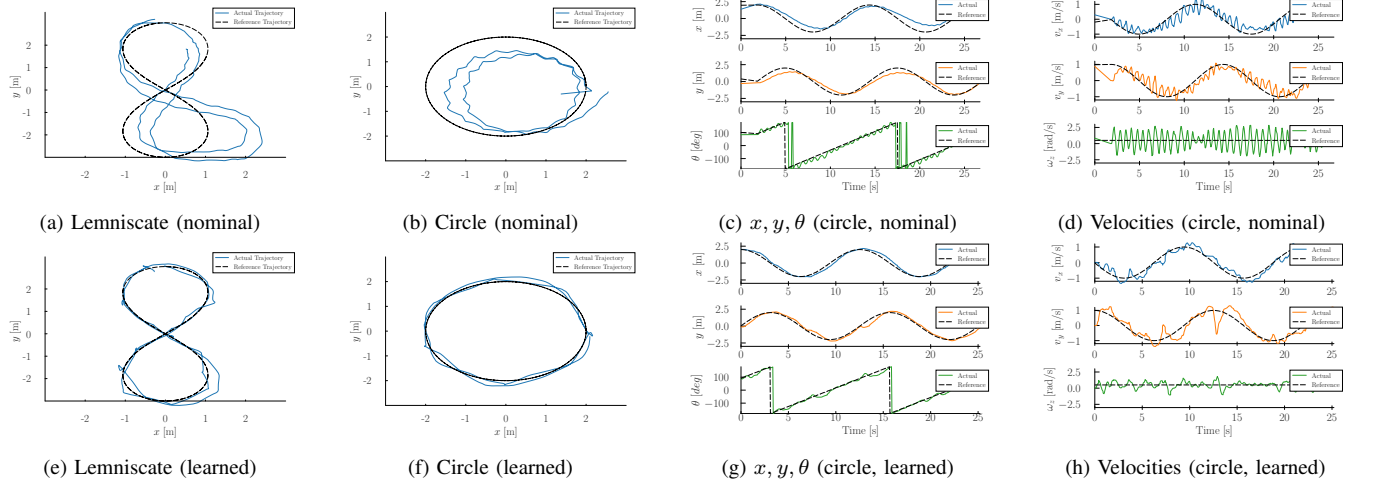


Fig. 5: Evaluating nominal and learned model performance in tracking circular and lemniscate trajectories.

dynamics. The unknown time constant α is learned as part of the control matrix $\mathbf{g}_\theta(\mathbf{q})$ during the training process.

Given Hamiltonian network design, described in Sec. III, we further specify that both the generalized mass and dissipation matrices have a block-diagonal form: $\mathbf{M}_\theta(\mathbf{q}) = \text{diag}(\mathbf{M}_{1,\theta}(\mathbf{q}), \mathbf{M}_{2,\theta}(\mathbf{q}))$, $\mathbf{D}_\theta(\mathbf{q}) = \text{diag}(\mathbf{D}_{v,\theta}(\mathbf{q}), \mathbf{D}_{\omega,\theta}(\mathbf{q}))$ and only the first and last row of $\mathbf{g}_\theta(\mathbf{q})$ are non-zero and to be learned from data. This aligns with the nonholonomic constraint discussed in Sec. IV-B. The model was trained for 400 iterations with potential energy $\mathcal{V}_\theta(\mathbf{q}) = 0$ and a fixed nominal model with $\bar{\mathbf{L}}(\mathbf{q}) = \text{block-diag}[\sqrt{m}\mathbf{I} \quad \Gamma]^\top$, $\bar{\mathbf{A}}(\mathbf{q}) = 0.1\mathbf{I}$, $\mathbf{g}_0(\mathbf{q}) = \begin{bmatrix} \frac{\eta}{\alpha r} & 0 & 0 & 0 & 0 & -\frac{\eta w}{2\alpha r} \\ \frac{\eta}{\alpha r} & 0 & 0 & 0 & 0 & \frac{\eta w}{2\alpha r} \end{bmatrix}^\top$ where the nominal mass $m = 16\text{kg}$, inertia $\Gamma\Gamma^\top = 0.4\mathbf{I}$, wheel radius $r = 0.098\text{m}$, vehicle width $w = 0.31\text{m}$, wheel inertia $\eta = 0.0458\text{kgm}^2$ and time constant $\alpha = 0.147$ [56].

Fig. 4 presents both training and testing results. The generalized mass and input gain converged similarly to the simulation results in Sec. V-B, while the generalized inertia has off-diagonal elements due to the LiDAR's off-center po-

sition. The dissipation matrices have large values in $[\mathbf{D}_{v,\theta}]_{1,1}$ and $[\mathbf{D}_{\omega,\theta}]_{1,1}$ consistent with the nonholonomic constraints, while $[\mathbf{D}_{v,\theta}]_{0,0}$ and $[\mathbf{D}_{\omega,\theta}]_{2,2}$ show smaller friction along x axis and around z axis.

Finally, we demonstrate our energy-based controller's performance using both nominal and learned models, in two experiments: (i) pose stabilization and (ii) trajectory tracking. For pose stabilization, both models successfully achieved the desired pose, as shown in Fig. 4, utilizing gains we tuned for the nominal model with $(k_p, k_{R_1}, k_{R_2}, \mathbf{K}_d) = (0.5, 0.5, 0.5, \text{block-diag}(1.5\mathbf{I}, 0.05\mathbf{I}))$. But in case of tracking, the learned model outperforms the nominal model as seen in Figure 5 without any additional controller tuning.

VI. CONCLUSIONS

We developed an approach for learning robot dynamics from point-cloud observations via Hamiltonian neural ODEs, design an energy-shaping trajectory tracking controller for the learned dynamics and apply to wheeled nonholonomic robots. Future work will focus on dynamics learning from images and online adaptation to changing conditions.

REFERENCES

- [1] I. Markovsky and F. Dörfler, “Behavioral systems theory in data-driven analysis, signal processing, and control,” *Annual Reviews in Control*, vol. 52, pp. 42–64, 2021.
- [2] T. M. Moerland, J. Broekens, A. Plaat, and C. M. Jonker. Now Foundations and Trends, 2023.
- [3] L. Ljung, “System identification,” *Wiley encyclopedia of electrical and electronics engineering*, 1999.
- [4] D. Nguyen-Tuong and J. Peters, “Model learning for robot control: a survey,” *Cognitive processing*, vol. 12, no. 4, 2011.
- [5] M. Deisenroth and C. E. Rasmussen, “PILCO: A model-based and data-efficient approach to policy search,” in *International Conference on Machine Learning (ICML)*, 2011.
- [6] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, “Information theoretic MPC for model-based reinforcement learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [7] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Multistep neural networks for data-driven discovery of nonlinear dynamical systems,” *arXiv preprint:1801.01236*, 2018.
- [8] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep reinforcement learning in a handful of trials using probabilistic dynamics models,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [9] M. A. Roehrl, T. A. Runkler, V. Brandstetter, M. Tokic, and S. Obermayer, “Modeling system dynamics with physics-informed neural networks based on Lagrangian mechanics,” *arXiv preprint:2005.14617*, 2020.
- [10] M. Lutter, C. Ritter, and J. Peters, “Deep Lagrangian networks: using physics as model prior for deep learning,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [11] J. K. Gupta, K. Menda, Z. Manchester, and M. J. Kochenderfer, “A general framework for structured learning of mechanical systems,” *arXiv preprint:1902.08705*, 2019.
- [12] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, “Lagrangian neural networks,” in *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.
- [13] M. Lutter, K. Listmann, and J. Peters, “Deep Lagrangian Networks for end-to-end learning of energy-based control for under-actuated systems,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [14] S. Greydanus, M. Dzamba, and J. Yosinski, “Hamiltonian neural networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [15] T. Bertalan, F. Dietrich, I. Mezić, and I. G. Kevrekidis, “On learning Hamiltonian systems from data,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, no. 12, 2019.
- [16] Z. Chen, J. Zhang, M. Arjovsky, and L. Bottou, “Symplectic recurrent neural networks,” *International Conference on Learning Representations (ICLR)*, 2020.
- [17] M. Finzi, K. A. Wang, and A. G. Wilson, “Simplifying Hamiltonian and Lagrangian neural networks via explicit constraints,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020.
- [18] Y. D. Zhong, B. Dey, and A. Chakraborty, “Symplectic ODE-Net: learning Hamiltonian dynamics with control,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [19] J. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar, “Integrating physics-based modeling with machine learning: A survey,” *arXiv preprint:2003.04919*, 2020.
- [20] T. Duong and N. Atanasov, “Hamiltonian-based Neural ODE Networks on the SE(3) Manifold For Dynamics Learning and Control,” in *Robotics: Science and Systems (RSS)*, Virtual, July 2021.
- [21] A. Havens and G. Chowdhary, “Forced variational integrator networks for prediction and control of mechanical systems,” in *Learning for Dynamics and Control (LADC)*, 2021, pp. 1142–1153.
- [22] V. Duruisseaux, T. Duong, M. Leok, and N. Atanasov, “Lie Group Forced Variational Integrator Networks for Learning and Control of Robot Systems,” in *Learning for Dynamics and Control (LADC)*, 2023, pp. 731–744.
- [23] C. D. McKinnon and A. P. Schoellig, “Meta learning with paired forward and inverse models for efficient receding horizon control,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3240–3247, 2021.
- [24] T. Duong and N. Atanasov, “Adaptive control of se (3) hamiltonian dynamics with learned disturbance features,” *IEEE Control Systems Letters*, vol. 6, pp. 2773–2778, 2022.
- [25] D. D. Fan, J. Nguyen, R. Thakker, N. Alatur, A.-a. Agha-mohammadi, and E. A. Theodorou, “Bayesian learning-based adaptive control for safety critical systems,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4093–4099.
- [26] J. Harrison, A. Sharma, R. Calandra, and M. Pavone, “Control adaptation via meta-learning dynamics,” in *Workshop on Meta-Learning at NeurIPS*, vol. 2018, 2018.
- [27] S. M. Richards, N. Azizan, J.-J. Slotine, and M. Pavone, “Adaptive-Control-Oriented Meta-Learning for Nonlinear Systems,” in *Robotics: Science and Systems*, Virtual, July 2021.
- [28] G. Shi, K. Azizzadenesheli, S.-J. Chung, and Y. Yue, “Meta-adaptive nonlinear control: Theory and algorithms,” *arXiv:2106.06098*, 2021.
- [29] M. O’Connell, G. Shi, X. Shi, and S.-J. Chung, “Meta-learning-based robust adaptive flight control under uncertain wind conditions,” *arXiv preprint:2103.01932*, 2021.
- [30] S. Chen and J. T. Wen, “Neural-learning trajectory tracking control of flexible-joint robot manipulators with unknown dynamics,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 128–135.
- [31] A. Romero, S. Govil, G. Yilmaz, Y. Song, and D. Scaramuzza, “Weighted maximum likelihood for controller tuning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1334–1341.
- [32] A. Zhang, R. T. McAllister, R. Calandra, Y. Gal, and S. Levine, “Learning invariant representations for reinforcement learning without reconstruction,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [33] D. Kamran, J. Zhu, and M. Lauer, “Learning path tracking for real car-like mobile robots from simulation,” in *European Conference on Mobile Robots (ECMR)*, 2019.
- [34] Y. Tian, K. Zhang, R. Tedrake, and S. Sra, “Can direct latent model learning solve linear quadratic gaussian control?” in *Learning for Dynamics and Control Conference (LADC)*, 2023, pp. 51–63.
- [35] T. Zhou, Y. Jae Lee, S. X. Yu, and A. A. Efros, “FlowWeb: Joint Image Set Alignment by Weaving Consistent, Pixel-Wise Correspondences,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [36] T. Zhou, P. Krahenbuhl, M. Aubry, Q. Huang, and A. A. Efros, “Learning dense correspondence via 3d-guided cycle consistency,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 117–126.
- [37] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *IEEE/CVF International Conference on Computer Vision (CVPR)*, 2017, pp. 2223–2232.
- [38] C. Zach, M. Klopschitz, and M. Pollefeys, “Disambiguating visual relations using loop constraints,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1426–1433.
- [39] K. Wilson and N. Snavely, “Network principles for sfm: Disambiguating repeated structures with local context,” in *IEEE International Conference on Computer Vision (ICCV)*, December 2013.
- [40] D. Dwibedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman, “Temporal cycle-consistency learning,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [41] H. Wu and X. Wang, “Contrastive learning of image representations with cross-video cycle-consistency,” in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 10 129–10 139.
- [42] Q. Zhang, T. Xiao, A. A. Efros, L. Pinto, and X. Wang, “Learning cross-domain correspondence for control with dynamics cycle-consistency,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [43] T. Duong, A. Altawaitan, J. Stanley, and N. Atanasov, “Port-Hamiltonian Neural ODE Networks on Lie Groups For Robot Dynamics Learning and Control,” *arXiv preprint:2401.09520*, 2024.
- [44] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, “Neural ordinary differential equations,” *Advances in neural information processing systems (NeurIPS)*, vol. 31, 2018.
- [45] Q.-Y. Zhou, J. Park, and V. Koltun, “Fast global registration,” in *European Conference on Computer Vision (ECCV)*, 2016, pp. 766–782.
- [46] A. Van Der Schaft and D. Jeltsema, “Port-Hamiltonian systems theory:

- An introductory overview,” *Foundations and Trends in Systems and Control*, vol. 1, no. 2-3, 2014.
- [47] Z. Wang and P. Goldsmith, “Modified energy-balancing-based control for the tracking problem,” *IET Control Theory & Applications*, vol. 2, no. 4, 2008.
- [48] C. Souza, G. V. Raffo, and E. B. Castelan, “Passivity based control of a quadrotor UAV,” *IFAC Proceedings Volumes*, vol. 47, no. 3, 2014.
- [49] M. Aicardi, G. Casalino, A. Bicchi, and A. Balestrino, “Closed loop steering of unicycle like vehicles via Lyapunov techniques,” *IEEE Robotics and Automation Magazine (RAM)*, vol. 2, no. 1, pp. 27–35, 1995.
- [50] W. F. Lages, *A Backstepping Non-smooth Controller for ROS-Based Differential-Drive Mobile Robots*. Cham: Springer International Publishing, 2019, pp. 275–314. [Online]. Available: https://doi.org/10.1007/978-3-319-91590-6_9
- [51] R. Hartley, J. Trumpf, Y. Dai, and H. Li, “Rotation averaging,” *International Journal of Computer Vision*, vol. 103, pp. 267–305, 2013.
- [52] A. Altwaitan, J. Stanley, S. Ghosal, T. Duong, and N. Atanasov, “Hamiltonian Dynamics Learning from Point Cloud Observations for Nonholonomic Mobile Robot Control,” University of California, San Diego, Tech. Rep., 2023. [Online]. Available: <https://altwaitan.github.io/DLFO/>
- [53] S. Boyd and S. Sastry, “Necessary and sufficient conditions for parameter convergence in adaptive control,” *Automatica*, vol. 22, no. 6, pp. 629–639, 1986.
- [54] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, “SAPIEN: A simulated part-based interactive environment,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [55] M. Faessler, D. Falanga, and D. Scaramuzza, “Thrust mixing, saturation, and body-rate control for accurate aggressive quadrotor flight,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 476–482, 2017.
- [56] W. Yu, O. Chuy, E. G. Collins, and P. Hollis, “Dynamic modeling of a skid-steered wheeled vehicle with experimental verification,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009, pp. 4212–4219.