Safe Autonomous Navigation for Systems with Learned SE(3) Hamiltonian Dynamics

Zhichao Li[*]   ZHICHAOLI@UCSD.EDU
Thai Duong[*]   TDUONG@UCSD.EDU
Nikolay Atanasov   NATANASOV@UCSD.EDU
Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093

Abstract

Safe autonomous navigation in unknown environments is an important problem for mobile robots. This paper proposes techniques to learn the dynamics model of a mobile robot from trajectory data and synthesize a tracking controller with safety and stability guarantees. The state of a rigid-body robot usually contains its position, orientation, and generalized velocity and satisfies Hamilton's equations of motion. Instead of a hand-derived dynamics model, we use a dataset of state-control trajectories to train a translation-equivariant nonlinear Hamiltonian model represented as a neural ordinary differential equation (ODE) network. The learned Hamiltonian model is used to synthesize an energy-shaping passivity-based controller and derive conditions which guarantee safe regulation to a desired reference pose. We enable adaptive tracking of a desired path, subject to safety constraints obtained from obstacle distance measurements. The trade-off between the robot's energy and the distance to safety constraint violation is used to adaptively govern a reference pose along the desired path. Our safe adaptive controller is demonstrated on a simulated hexarotor robot navigating in an unknown environments.

Keywords: dynamics learning, neural ODE network, reference governor, safe tracking control

1. Introduction

Designing controllers that handle safety constraints and guarantee stability is an important problem in safety-critical applications of robotics, such as autonomous driving (Ames et al., 2014b; Shalev-Shwartz et al., 2016), locomotion (Ames et al., 2014a) or medical robotics (Yip and Camarillo, 2014). Safety depends on the system states, governed by the system dynamics, and the environment constraints. This leads to two requirements for designing provably safe controllers: 1) the availability of an accurate dynamics model and 2) the satisfaction of time-varying safety constraints that are only known at runtime.

The first requirement has motivated machine learning techniques for system dynamics learning, e.g. based on Gaussian processes (Deisenroth et al., 2015; Kabzan et al., 2019) or neural networks (Raissi et al., 2018; Chua et al., 2018). For physical systems, recent works (Lutter et al., 2019; Zhong et al., 2019; Duong and Atanasov, 2021b) design the model architecture to encode a Lagrangian or Hamiltonian formulation of robot dynamics (Lurie, 2013; Holm, 2008), which a black-box model might struggle to infer. Zhong et al. (2019) use a differentiable neural ODE solver (Chen et al., 2018) to generate predicted state trajectories in a Hamiltonian formulation. A loss function is back-

---

[*] These authors contributed equally.

propagated through the ODE solver to update the model parameters. Duong and Atanasov (2021b) extend this approach by imposing both Hamiltonian dynamics and $SE(3)$ pose constraints on the ODE structure. A Hamiltonian-based model architecture also simplifies the design of stable regulation or tracking control by energy shaping (Zhong et al., 2019; Duong and Atanasov, 2021a,b). The key idea of energy-based shaping, known as interconnection and damping assignment passivity-based control (IDA-PBC) (Van Der Schaft and Jeltsema, 2014), is to inject additional energy via the control input into the system to achieve a desired total energy, minimized at a desired set point.

The second requirement related to safety guarantees has gained significant attention in planning and control. Model predictive control (MPC) methods (Borrelli et al., 2017; Grüne and Pannek, 2017; Bravo et al., 2006; Mayne et al., 2000) include safety constraints in an optimization problem, which is typically solved by discretizing time with linearized dynamics. Reachability-based techniques (Herbert et al., 2017; Majumdar and Tedrake, 2017; Kousik et al., 2020) work directly for nonlinear systems and offer strong safety guarantees but have high computation cost and scalability issues for high-dimensional systems. Control barrier functions (CBFs) with quadratic programming (QP) (Ames et al., 2014b, 2017, 2019) offer an elegant and efficient framework for real-time safe control synthesis. However, construct a valid CBF (Ames et al., 2019) that guarantees the feasibility of the QP problem (Xu, 2018) at all times is challenging. Given a stabilizing regulation controller, reference governor techniques (Bemporad, 1998; Kolmanovsky et al., 2014; Garone and Nicotra, 2016) maintain a virtual governor system to adaptively generate a regulation so that the system follows reference commands safely. Recent work (Arslan and Koditschek, 2017; Li et al., 2020) achieves safe navigation in unknown environments but is limited to feedback-linearizable systems.

In this paper, we consider both requirements for rigid-body robot systems, whose states are described by their $SE(3)$ pose and generalized velocity. We assume that the robot dynamics are unknown but, as a physical system, satisfy Hamilton's equations of motion over the $SE(3)$ manifold. We consider a training set of state-control trajectories, from past experiments or collected by a human operator, and seek to safely track a desired position path with safety constraints obtained online from distance-to-obstacles measurements. We learn a $SE(3)$ Hamiltonian model of the system dynamics using a neural ODE network (Duong and Atanasov, 2021b). As the robot dynamics are equivariant to translation, we offset the trajectories to start from the origin and train a translation-equivariant Hamiltonian neural ODE model. The Hamiltonian structure of the learned model offers an energy-based regulation controller with the total energy of the system viewed as a Lyapunov function. This, in turn, enables us to enforce safety constraints using reference governor techniques without the need to linearize the system dynamics. Inspired by constraint embedding techniques (Garone and Nicotra, 2016), we impose safety constraints, based on the sensor measurements, on the Lyapunov function. We use the trade-off between the distance from constraint violation and the system energy level to regulate a reference governor and achieve safe and stable position tracking in an unknown environment.

Contributions. In summary, the contributions of this paper are 1) a translation-equivariant $SE(3)$ Hamiltonian dynamics learning approach and 2) a tracking control design for $SE(3)$ Hamiltonian systems with stability and safety guarantees. Our dynamics learning and tracking control techniques are demonstrated on a simulated hexarotor robot using a depth sensor to navigate in unknown environments.

2. Problem Statement

Consider a robot modeled as a rigid body with position $\mathbf{p} \in \mathbb{R}^3$, orientation $\mathbf{R} \in SO(3)$, body-frame linear velocity $\mathbf{v} \in \mathbb{R}^3$, and body-frame angular velocity $\boldsymbol{\omega} \in \mathbb{R}^3$. Let $\mathfrak{q} = [\mathbf{p}^\top \ \mathbf{r}_1^\top \ \mathbf{r}_2^\top \ \mathbf{r}_3^\top]^\top \in \mathbb{R}^{12}$ denote the robot's generalized coordinates, where $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3 \in \mathbb{R}^3$ are the rows of the rotation matrix $\mathbf{R}$. Let $\boldsymbol{\zeta} = [\mathbf{v}^\top \ \boldsymbol{\omega}^\top]^\top \in \mathbb{R}^6$ denote the robot's generalized velocity. The generalized momentum $\mathfrak{p}$ of the system is defined as:

$$\mathfrak{p} = \mathbf{M}(\mathfrak{q})\boldsymbol{\zeta} \in \mathbb{R}^6, \tag{1}$$

where $\mathbf{M}(\mathfrak{q}) \succ 0$ denotes a positive-definite 6×6 generalized mass matrix. Let $\mathbf{x} = (\mathfrak{q}, \mathfrak{p}) \in \mathbb{R}^{18}$ be the robot state. The Hamiltonian, $\mathcal{H}(\mathfrak{q}, \mathfrak{p})$, captures the total energy of the system as the sum of the kinetic energy $\mathcal{T}(\mathfrak{q}, \mathfrak{p}) = \frac{1}{2}\mathfrak{p}^\top \mathbf{M}(\mathfrak{q})^{-1}\mathfrak{p}$ and the potential energy $\mathcal{U}(\mathfrak{q})$:

$$\mathcal{H}(\mathfrak{q}, \mathfrak{p}) = \mathcal{T}(\mathfrak{q}, \mathfrak{p}) + \mathcal{U}(\mathfrak{q}) = \frac{1}{2}\mathfrak{p}^\top \mathbf{M}(\mathfrak{q})^{-1}\mathfrak{p} + \mathcal{U}(\mathfrak{q}). \tag{2}$$

As a mechanical system, the time evolution of the state $\mathbf{x}$ is governed by Hamilton's equations of motion (Lee et al., 2017; Duong and Atanasov, 2021b):

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\mathbf{u}, \quad \mathbf{x}(t_0) = \mathbf{x}_0, \tag{3}$$

where $\mathbf{u} \in \mathbb{R}^p$ is the control input, $\mathbf{f}(\mathbf{x}) = \begin{bmatrix} \mathbf{0} & \mathfrak{q}^\times \\ -\mathfrak{q}^{\times\top} & \mathfrak{p}^\times \end{bmatrix} \begin{bmatrix} \nabla_\mathfrak{q}\mathcal{H} \\ \nabla_\mathfrak{p}\mathcal{H} \end{bmatrix}$, $\mathbf{G}(\mathbf{x}) = \begin{bmatrix} \mathbf{0} \\ \mathbf{B}(\mathfrak{q}) \end{bmatrix}$, and $\mathbf{B}(\mathfrak{q}) \in \mathbb{R}^{6 \times p}$ is an input gain matrix. The operators $\mathfrak{q}^\times$, $\mathfrak{p}^\times$, and the hat map $\hat{\mathbf{w}}$ for $\mathbf{w} \in \mathbb{R}^3$ are defined as:

$$\mathfrak{q}^\times = \begin{bmatrix} \mathbf{R}^\top & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{r}}_1^\top & \hat{\mathbf{r}}_2^\top & \hat{\mathbf{r}}_3^\top \end{bmatrix}^\top, \quad \mathfrak{p}^\times = \begin{bmatrix} \mathfrak{p}_\mathbf{v} \\ \mathfrak{p}_{\boldsymbol{\omega}} \end{bmatrix}^\times = \begin{bmatrix} \mathbf{0} & \hat{\mathfrak{p}}_\mathbf{v} \\ \hat{\mathfrak{p}}_\mathbf{v} & \hat{\mathfrak{p}}_{\boldsymbol{\omega}} \end{bmatrix}, \quad \hat{\mathbf{w}} = \begin{bmatrix} 0 & -w_3 & w_2 \\ w_3 & 0 & -w_1 \\ -w_2 & w_1 & 0 \end{bmatrix}.$$

We consider the case that the parameters of the Hamiltonian dynamics model in (3), including the mass $\mathbf{M}(\mathfrak{q})$, potential energy $\mathcal{U}(\mathfrak{q})$, and input matrix $\mathbf{B}(\mathfrak{q})$, are unknown. Instead, we are given a trajectory dataset $\mathcal{D} = \{t_{0:N}^{(i)}, \mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)}, \mathbf{u}^{(i)}\}_{i=1}^D$ consisting of D sequences of generalized coordinates and velocities $(\mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)})$ at times $t_0^{(i)} < t_1^{(i)} < \ldots < t_N^{(i)}$, collected by applying a constant control input $\mathbf{u}^{(i)}$ to the system with initial condition $(\mathfrak{q}_0^{(i)}, \boldsymbol{\zeta}_0^{(i)})$. We aim to learn the dynamics from the data set $\mathcal{D}$ and design a control policy $\mathbf{u} = \boldsymbol{\pi}(\mathbf{x})$ such that the robot follows a desired reference path without violating safety constraints in an unknown environment. Let $\mathcal{O} \subset \mathbb{R}^3$ and $\mathcal{F} := \mathbb{R}^3 \setminus \mathcal{O}$ denote the unsafe (e.g., obstacle) set and the safe (obstacle-free) set, respectively. Denote the interior of $\mathcal{F}$ as $\text{int}(\mathcal{F})$. We assume that $\mathcal{O}$ is not known a priori but the robot can sense the distance $\bar{d}(\mathbf{p}, \mathcal{O})$ from its position $\mathbf{p}$ to $\mathcal{O}$ locally with a limited sensing range $\beta > 0$:

$$\bar{d}(\mathbf{p}, \mathcal{O}) := \min\{d(\mathbf{p}, \mathcal{O}), \beta\}, \tag{4}$$

where $d(\mathbf{p}, \mathcal{O}) := \inf_{\mathbf{a} \in \mathcal{O}} \|\mathbf{p} - \mathbf{a}\|$ denotes the Euclidean distance from $\mathbf{p}$ to the set $\mathcal{O}$. The safe autonomous navigation problem considered in this paper is summarized below.

Problem 1 *Let $\mathcal{D} = \{t_{0:N}^{(i)}, \mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)}, \mathbf{u}^{(i)}\}_{i=1}^D$ be a training dataset of state-control trajectories obtained from a robot with unknown Hamiltonian dynamics in (3). Let $\mathbf{r} : [0, 1] \mapsto Int(\mathcal{F})$ be*

a continuous function specifying a desired position reference path for the robot. Assume that the reference path starts at the initial robot position at time t_0, i.e., $\mathbf{r}(0) = \mathbf{p}(t_0) \in Int\,(\mathcal{F})$. Using local distance observations $\bar{d}(\mathbf{p}(t), \mathcal{O})$ of the unsafe set $\mathcal{O}$ in an unknown environment, design a control policy $\boldsymbol{\pi}(\mathbf{x})$ so that the position $\mathbf{p}(t)$ of the closed-loop system converges asymptotically to $\mathbf{r}(1)$, while remaining safe, i.e., $\mathbf{p}(t) \in \mathcal{F}$ for all $t \geq t_0$.

3. Learning Hamiltonian Dynamics on the $SE(3)$ Manifold

3.1. Training a translation-equivariant $SE(3)$ Hamiltonian dynamics model

The training dataset $\mathcal{D} = \{t_{0:N}^{(i)}, \mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)}, \mathbf{u}^{(i)}\}_{i=1}^{D}$ can be obatined by measuring the generalized coordinates and velocities of the system at times $t_0^{(i)} < t_1^{(i)} < \ldots < t_N^{(i)}$ using an odometry algorithm (Delmerico and Scaramuzza, 2018; Mohamed et al., 2019) or a motion capture system. The control input $\mathbf{u}^{(i)}$ can be generated by manually driving the robot or using an existing controller. Since the system dynamics do not change if we shift the position $\mathbf{p}$ to any points in the world frame, we offset the trajectories in the dataset $\mathcal{D}$ so that they start from the position $\mathbf{0}$ and learn the system dynamics around the origin. This is sufficient for control purposes, e.g. using the control design in Sec. 4.1, the input driving the system from state $\mathbf{x}$ with position $\mathbf{p}$ to desired state $\mathbf{x}^*$ with position $\mathbf{p}^*$ is the same as the one driving the system from state $\mathbf{x}$ with position $\mathbf{0}$ to the desired state $\mathbf{x}^*$ with the offset position $\mathbf{p}^* - \mathbf{p}$.

Since the momentum $\mathfrak{p}$ is not directly available from the dataset $\mathcal{D}$, we use the time derivative of the generalized velocity, derived from (1):

$$\dot{\boldsymbol{\zeta}} = \left(\frac{d}{dt}\mathbf{M}^{-1}(\mathfrak{q})\right)\mathfrak{p} + \mathbf{M}^{-1}(\mathfrak{q})\dot{\mathfrak{p}}. \tag{5}$$

Eq. (3) and (5) describe the Hamiltonian dynamics with unknown inverse generalized mass matrix $\mathbf{M}(\mathfrak{q})^{-1}$, input matrix $\mathbf{B}(\mathfrak{q})$, and potential energy $\mathcal{U}(\mathfrak{q})$, which we aim to approximate by three neural networks $\mathbf{M}_{\boldsymbol{\theta}}(\mathfrak{q})^{-1}, \mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$ and $\mathcal{U}_{\boldsymbol{\theta}}(\mathfrak{q})$, respectively, with parameters $\boldsymbol{\theta}$.

To optimize the parameters $\boldsymbol{\theta}$, we use a neural ODE framework that encodes the Hamiltonian dynamics (3) and (5) with $\mathbf{M}_{\boldsymbol{\theta}}(\mathfrak{q}), \mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$ and $\mathcal{U}_{\boldsymbol{\theta}}(\mathfrak{q})$ in the network structure (Fig. 1(a)). The forward pass rolls out the Hamiltonian dynamics (3) and (5) with the neural networks $\mathbf{M}_{\boldsymbol{\theta}}(\mathfrak{q}), \mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$ and $\mathcal{U}_{\boldsymbol{\theta}}(\mathfrak{q})$ using a neural ODE solver (Chen et al., 2018) to obtain a predicted sequence $(\bar{\mathfrak{q}}_{0:N}^{(i)}, \bar{\boldsymbol{\zeta}}_{0:N}^{(i)})$ at times $t_0^{(i)} < t_1^{(i)} < \ldots < t_N^{(i)}$ for each $i = 1, \ldots, D$. The loss function is defined as $\mathcal{L} = \sum_{i=1}^{D} \sum_{n=1}^{N} c(\mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)}, \bar{\mathfrak{q}}_{0:N}^{(i)}, \bar{\boldsymbol{\zeta}}_{0:N}^{(i)})$, where the distance metric c is defined as the sum of position, orientation, and velocity errors:

$$c\left(\mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)}, \bar{\mathfrak{q}}_{0:N}^{(i)}, \bar{\boldsymbol{\zeta}}_{0:N}^{(i)}\right) = c_{\mathbf{P}}(\mathbf{p}, \bar{\mathbf{p}}) + c_{\mathbf{R}}(\mathbf{R}, \bar{\mathbf{R}}) + c_{\boldsymbol{\zeta}}(\boldsymbol{\zeta}, \bar{\boldsymbol{\zeta}}), \tag{6}$$

where $c_{\mathbf{P}}(\mathbf{p}, \bar{\mathbf{p}}) = \|\mathbf{p} - \bar{\mathbf{p}}\|_2^2$, $c_{\boldsymbol{\zeta}}(\boldsymbol{\zeta}, \bar{\boldsymbol{\zeta}}) = \|\boldsymbol{\zeta} - \bar{\boldsymbol{\zeta}}\|_2^2$, and $c_{\mathbf{R}}(\mathbf{R}, \bar{\mathbf{R}}) = \|\left(\log(\bar{\mathbf{R}}\mathbf{R}^{\top})\right)^{\vee}\|_2^2$. The log-map $\log(\cdot) : SE(3) \mapsto \mathfrak{so}(3)$ is the inverse of the exponential map, returning a skew-symmetric matrix in $\mathfrak{so}(3)$ from a rotation matrix in $SE(3)$, and the $\vee$-map $(\cdot)^{\vee} : \mathfrak{so}(3) \mapsto \mathbb{R}^3$ is the inverse of the hat map $\hat{(\cdot)}$ in Sec. 2. The parameters $\boldsymbol{\theta}$ are optimized via gradient descent by back-propagating the loss through the neural ODE solver. This is done using the adjoint method, where the gradient $\partial\mathcal{L}/\partial\boldsymbol{\theta}$ is calculated by a single call to a reverse-time ODE starting from $t = t_N$ at $(\mathfrak{q}_N^{(i)}, \boldsymbol{\zeta}_N^{(i)})$.
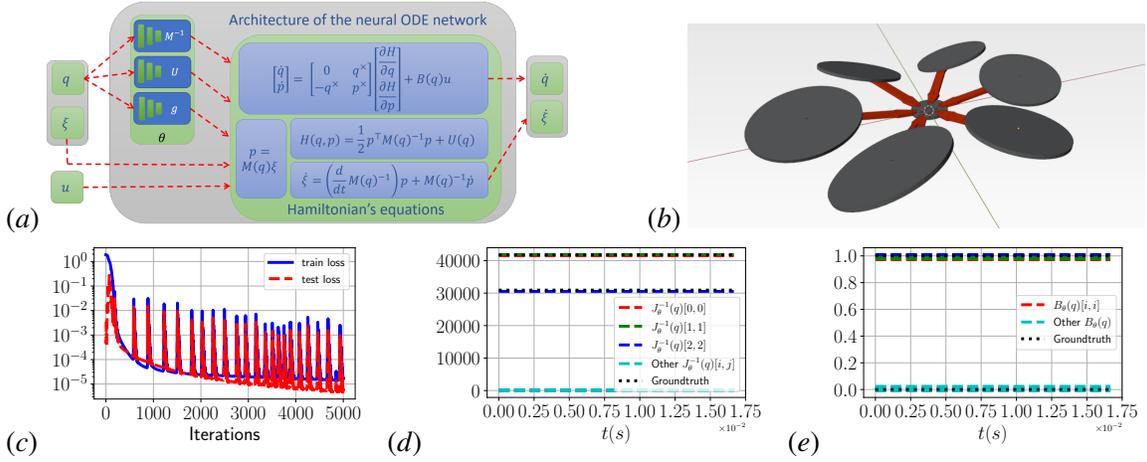
Figure 1: $SE(3)$ Hamiltonian neural ODE network: (a) network architecture, (b) simulated hexarotor for evaluation, (c) training loss, (d) learned inverse inertia $\mathbf{J}_{\boldsymbol{\theta}}(\mathfrak{q})^{-1}$, and (e) learned input matrix $\mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$ along a test trajectory, evaluated on the simulated hexarotor.

3.2. Evaluation of the Hamiltonian dynamics model of a simulated hexarotor

We consider a simulated hexarotor unmanned aerial vehicle (UAV) (Fig. 1(b)) with fixed-tilt rotors pointing in different directions (Rajappa et al. (2015)) modeled as a fully-actuated rigid body with mass $m = 0.027$ and inertia matrix $\mathbf{J} = 10^{-5}\text{diag}([2.4, 2.4, 3.2])$. The robot's ground-truth dynamics satisfy Hamilton's equations in (3) with generalized mass $\mathbf{M}(\mathfrak{q}) = \text{diag}(m\mathbf{I}, \mathbf{J})$, potential energy $\mathcal{U}(\mathfrak{q}) = mgz$, and the input matrix $\mathbf{B}(\mathfrak{q}) = \mathbf{I}$. The control input $\mathbf{u}$ is a 6-dimensional wrench, including a 3-dimensional force and a 3-dimensional torque. Since the mass m of the hexarotor can be easily measured, we assume the mass m is known, leading to a known potential energy $\mathcal{U}(\mathfrak{q}) = mg \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \mathbf{p}$, where $\mathbf{p}$ is the UAV position and $g \approx 9.8ms^{-2}$ is the gravitational acceleration. We approximate the inverse generalized mass matrix by $\mathbf{M}_{\boldsymbol{\theta}}(\mathfrak{q})^{-1} = \text{diag}(m^{-1}\mathbf{I}, \mathbf{J}_{\boldsymbol{\theta}}^{-1}(\mathfrak{q}))$ and learn the inverse inertia matrix $\mathbf{J}_{\boldsymbol{\theta}}(\mathfrak{q})^{-1}$ and the input matrix $\mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$ from data.

As the dynamics model is translation-invariant, we mimic manual flights in an area free of obstacles using a PID controller and drive the hexarotor from a random initial pose near the origin to a desired pose, generating 18 one-second trajectories. We shift the trajectories to start from the origin and create a dataset $\mathcal{D} = \{t_{0:N}^{(i)}, \mathfrak{q}_{0:N}^{(i)}, \boldsymbol{\zeta}_{0:N}^{(i)}, \mathbf{u}^{(i)})\}_{i=1}^{D}$ with $N = 5$ and $D = 432$. The Hamiltonian-based neural ODE network[1] is trained with the dataset $\mathcal{D}$, as described in Sec. 3, for 5000 iterations and learning rate 10^{-3}. Fig. 1(c) shows the loss function during training. Note that if we scale $\mathbf{M}_{\boldsymbol{\theta}}(\mathfrak{q})$ and the input matrix $\mathbf{B}(\mathfrak{q})$ by a constant γ, the dynamics of $(\mathfrak{q}, \boldsymbol{\zeta})$ in (3) and (5) do not change. Fig. 1(d) and 1(e) plot the scaled version of the learned inverse mass $\mathbf{J}_{\boldsymbol{\theta}}(\mathfrak{q})^{-1}$ and the input matrix $\mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$, converging to the constant ground truth values.

---

1. Code: https://thaipduong.github.io/SE3HamDL/

4. Safe Tracking using a Reference Governor

In this section, we first describe a passivity-based regulation controller for arbitrary pose stabilization in Sec. 4.1. We derive sufficient conditions for safety based on an invariant level set of the closed-loop system's Hamiltonian. In Sec. 4.2, we propose a reference governor control policy to adaptively generate a regulation pose along the desired path and achieve safe navigation.

4.1. Passivity-based control for learned Hamiltonian dynamics

Given the learned model of the system dynamics:

$$\dot{\mathbf{x}} = \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x}) + \mathbf{G}_{\boldsymbol{\theta}}(\mathbf{x})\mathbf{u}, \quad \mathbf{x}(t_0) = \mathbf{x}_0, \tag{7}$$

we want to find a control policy that stabilizes the system to a desired equilibrium $\mathbf{x}^* := (\mathfrak{q}^*, \mathbf{0})$ with desired generalized coordinates $\mathfrak{q}^* = (\mathbf{p}^*, \mathbf{R}^*)$ and zero momentum $\mathfrak{p}^* = \mathbf{0}$. We design a control policy $\mathbf{u} = \boldsymbol{\pi}(\mathbf{x}, \mathbf{x}^*)$ to shape the total energy (Hamiltonian) of the closed-loop system so that it achieves a minimum at the desired state $\mathbf{x}^* = (\mathfrak{q}^*, \mathbf{0})$. By injecting energy into the system through the controller $\mathbf{u} = \boldsymbol{\pi}(\mathbf{x}, \mathbf{x}^*)$, we aim to achieve the following desired Hamiltonian:

$$\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*) = \frac{1}{2}k_{\mathbf{p}}(\mathbf{p} - \mathbf{p}^*)^{\top}(\mathbf{p} - \mathbf{p}^*) + \frac{1}{2}k_{\mathbf{R}}\operatorname{tr}(\mathbf{I} - \mathbf{R}^{*\top}\mathbf{R}) + \frac{1}{2}(\mathfrak{p} - \mathfrak{p}^*)^{\top}\mathbf{M}_{\boldsymbol{\theta}}^{-1}(\mathfrak{q})(\mathfrak{p} - \mathfrak{p}^*), \tag{8}$$

where $k_{\mathbf{p}}$ and $k_{\mathbf{R}}$ are positive gains. We use interconnection and damping assignment passivity-based control (IDA-PBC) (Van Der Schaft and Jeltsema, 2014):

$$\mathbf{u} = \mathbf{B}_{\boldsymbol{\theta}}^{\dagger}(\mathfrak{q})\left(\mathfrak{q}^{\times\top}\nabla_{\mathfrak{q}}\mathcal{U}_{\boldsymbol{\theta}}(\mathfrak{q}) - \mathfrak{p}^{\times}\mathbf{M}_{\boldsymbol{\theta}}^{-1}(\mathfrak{q})\mathfrak{p} - \mathbf{e}(\mathfrak{q}, \mathfrak{q}^*)\right) + \left(-\mathbf{B}_{\boldsymbol{\theta}}^{\dagger}(\mathfrak{q})\mathbf{K}_{\mathbf{d}}\mathbf{M}_{\boldsymbol{\theta}}^{-1}(\mathfrak{q})\mathfrak{p}\right), \tag{9}$$

where $\mathbf{B}_{\boldsymbol{\theta}}^{\dagger}(\mathfrak{q}) = \left(\mathbf{B}_{\boldsymbol{\theta}}^{\top}(\mathfrak{q})\mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})\right)^{-1}\mathbf{B}_{\boldsymbol{\theta}}^{\top}(\mathfrak{q})$ is the pseudo-inverse of $\mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$, $\mathbf{K}_{\mathbf{d}} = \operatorname{diag}(k_{\mathbf{v}}\mathbf{I}, k_{\boldsymbol{\omega}}\mathbf{I})$ is a damping gain with positive terms $k_{\mathbf{v}}$, $k_{\boldsymbol{\omega}}$, and $\mathbf{e}(\mathfrak{q}, \mathfrak{q}^*)$ is the error between $\mathfrak{q}$ and $\mathfrak{q}^*$:

$$\mathbf{e}(\mathfrak{q}, \mathfrak{q}^*) = \begin{bmatrix}\mathbf{e}_{\mathbf{p}}(\mathfrak{q}, \mathfrak{q}^*)\\\mathbf{e}_{\mathbf{R}}(\mathfrak{q}, \mathfrak{q}^*)\end{bmatrix} = \begin{bmatrix}k_{\mathbf{p}}\mathbf{R}^{\top}(\mathbf{p} - \mathbf{p}^*)\\\frac{1}{2}k_{\mathbf{R}}\left(\mathbf{R}^{*\top}\mathbf{R} - \mathbf{R}^{\top}\mathbf{R}^*\right)^{\vee}\end{bmatrix}. \tag{10}$$

Lemma 1 *If the input gain matrix $\mathbf{B}_{\boldsymbol{\theta}}(\mathfrak{q})$ of the system in (7) is invertible, the control policy $\mathbf{u} = \boldsymbol{\pi}(\mathbf{x}, \mathbf{x}^*)$ in (9) always exists and asymptotically stabilizes the system to an arbitrary reference $\mathbf{x}^* = (\mathfrak{q}^*, \mathbf{0})$ with Lyapunov function given by the desired Hamiltonian $\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*)$ in (8).*

Proof Since $\dot{\mathcal{H}}_d(\mathbf{x}, \mathbf{x}^*) = -\mathfrak{p}^{\top}\mathbf{M}_{\boldsymbol{\theta}}^{-1}(\mathfrak{q})\mathbf{K}_d\mathbf{M}_{\boldsymbol{\theta}}^{-1}(\mathfrak{q})\mathfrak{p} \leq 0$, the proof is obtained by LaSalle's invariance principle (Khalil, 2002). See Appendix A in the extended version (Li et al., 2021). ∎

Next, given the closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x}) + \mathbf{G}_{\boldsymbol{\theta}}(\mathbf{x})\boldsymbol{\pi}(\mathbf{x}, \mathbf{x}^*), \quad \mathbf{x}(t_0) = \mathbf{x}_0, \tag{11}$$

we derive conditions on the initial state $\mathbf{x}_0$ under which the position $\mathbf{p}$ converges to $\mathbf{p}^*$ safely, remaining in the safe set $\mathcal{F}$. We first define a dynamic safety margin (DSM) $\Delta E(\mathbf{x}, \mathbf{x}^*)$ for the Hamiltonian dynamics (11):

$$\Delta E(\mathbf{x}, \mathbf{x}^*) := \bar{d}^2(\mathbf{p}^*, \mathcal{O}) - 2\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*)/k_{\mathbf{p}}, \tag{12}$$

where $\bar{d}^2(\mathbf{p}^*, \mathcal{O})$ is the truncated distance to the unsafe set $\mathcal{O}$ in (4). Given a fixed desired point $\mathbf{x}^*$, the DSM function measures the the trade-off between safety, measured by $\bar{d}^2(\mathbf{p}^*, \mathcal{O})$, and system energy, measured by $\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*)$ and allows us to find a positively forward invariant set $\mathcal{S}(\mathbf{x}, \mathbf{x}^*)$ such that for any $\mathbf{x}_0 \in \mathcal{S}(\mathbf{x}, \mathbf{x}^*)$, the position $\mathbf{p}$ converges to $\mathbf{p}^*$ while remaining in the safe set $\mathcal{F}$.

Proposition 2 *Controller* (9) *renders* $\mathcal{S}(\mathbf{x}, \mathbf{x}^*) \coloneqq \{\mathbf{x} \mid \Delta E(\mathbf{x}, \mathbf{x}^*) \geq 0\}$ *a positively forward invariant set for system* (11). *Furthermore, if* $\mathbf{x}_0 \in \mathcal{S}(\mathbf{x}, \mathbf{x}^*)$, *then the state* $\mathbf{p}(t)$ *converges to* $\mathbf{p}^*$ *asymptotically without collisions, i.e.,* $d(\mathbf{p}(t), \mathcal{O}) \geq 0$ *for all* $t \geq t_0$.

Proof From Lemma 1, we see that $\mathcal{S}(\mathbf{x}, \mathbf{x}^*) = \{\mathbf{x} \mid \Delta E(\mathbf{x}, \mathbf{x}^*) \geq 0\}$ is a forward invariant set induced by $\mathcal{H}(\mathbf{x}, \mathbf{x}^*)$. Note that $\frac{2}{k_{\mathbf{p}}}\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*) \geq \|(\mathbf{p} - \mathbf{p}^*)\|^2$, hence, $d^2(\mathbf{p}^*, \mathcal{O}) \geq \bar{d}^2(\mathbf{p}^*, \mathcal{O}) \geq \frac{2}{k_{\mathbf{p}}}\mathcal{H}_d(\mathbf{x}_0, \mathbf{x}^*) \geq \frac{2}{k_{\mathbf{p}}}\mathcal{H}_d(\mathbf{x}(t), \mathbf{x}^*) \geq d^2(\mathbf{p}^*, \mathbf{p}(t))$. See Appendix B (Li et al., 2021) for details. ∎

Using the results for a fixed $\mathbf{x}^*$ in this section, we develop a reference governor system to adaptive change $\mathbf{x}^*$ over time in Sec. 4.2 so that the robot can safely track a desired path.

4.2. Reference governor design

We introduce a virtual system, called a *reference governor* (Bemporad, 1998) to adaptively track the path $\mathbf{r}$ defined in Sec. 2 and provide a time-varying reference $\mathbf{x}^*(t)$ for the actual system in (11). The motion of the governor system needs to be regulated to balance the energy of the Hamiltonian system with the distance to the unsafe set $\mathcal{O}$, keeping the safety margin in (12) positive. We define the governor as a first-order linear time-invariant system with state $\mathbf{g}(t) \in \mathbb{R}^3$ and dynamics:

$$\dot{\mathbf{g}} = -k_g(\mathbf{g} - \mathbf{u_g}), \quad k_g > 0. \tag{13}$$

The governor input $\mathbf{u_g}$ will be chosen to move the governor system along the reference path without violating the safety condition $\Delta E(\mathbf{x}, \mathbf{x}^*) \geq 0$ obtained in Proposition 2. Define a *local safe set* $\mathcal{LS}(\mathbf{x}, \mathbf{g})$ as a region around the governor state $\mathbf{g}$ that does not violate safety:

$$\mathcal{LS}(\mathbf{x}, \mathbf{g}) \coloneqq \left\{\mathbf{q} \in \mathbb{R}^m \mid \|\mathbf{q} - \mathbf{g}\|^2 \leq (1 + \epsilon)^{-1}\Delta E(\mathbf{x}, \mathbf{x}^*)\right\},$$

where $\epsilon > 0$ is arbitrarily, ensuring that $\mathcal{LS}(\mathbf{x}, \mathbf{g}) \subseteq \text{int}(\mathcal{F})$. The size of the local safe set determines how fast the governor can move along the reference path without endangering safety.

Definition 3 *A* local projected goal *at system-governor state* $(\mathbf{x}, \mathbf{g})$ *is a point* $\bar{\mathbf{g}} \in \mathcal{LS}(\mathbf{x}, \mathbf{g})$ *that is furthest along the reference path* $\mathbf{r}$:

$$\bar{\mathbf{g}} = \mathbf{r}(\sigma^*), \quad \sigma^* = \underset{\sigma \in [0,1]}{\text{argmax}}\{\sigma \mid \mathbf{r}(\sigma) \in \mathcal{LS}(\mathbf{x}, \mathbf{g})\}. \tag{14}$$

Choosing the governor input as $\mathbf{u_g} = \bar{\mathbf{g}}$ forces the governor to track the reference path adaptively, taking the safety condition $\Delta E(\mathbf{x}, \mathbf{x}^*) \geq 0$ into account. Given the local projected goal $\bar{\mathbf{g}}(t) \in \mathbb{R}^3$ and the governor state $\mathbf{g}(t) \in \mathbb{R}^3$, we also generate a desired reference state $\mathbf{x}^*(t)$ for the system in (11) by lifting $\mathbf{g}(t)$ to $\mathbb{R}^{18}$. We may choose $\mathbf{g}(t)$ as the desired position with zero desired velocity but, to provide guidance on the $SE(3)$ manifold, we need to also generate a desired orientation $\mathbf{R}^*(t)$. We construct a lifting function $\boldsymbol{\ell} : \mathcal{F} \times \text{int}(\mathcal{F}) \mapsto \mathbb{R}^{18}$ to obtain $\mathbf{x}^* = \boldsymbol{\ell}(\mathbf{g}, \bar{\mathbf{g}})$ as:

$$\boldsymbol{\ell}(\mathbf{g}, \bar{\mathbf{g}}) = \begin{bmatrix} \mathbf{p}^{*\top} & \mathbf{r}_1^{*\top} & \mathbf{r}_2^{*\top} & \mathbf{r}_3^{*\top}, \mathbf{0}^\top, \mathbf{0}^\top \end{bmatrix}^\top, \tag{15}$$
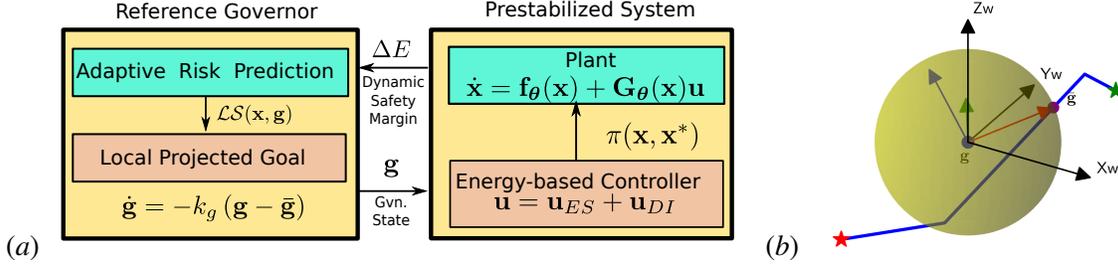
Figure 2: (a) Structure of the reference-governor tracking controller: a reference governor with state $\mathbf{g}$ adaptively tracks a point $\bar{\mathbf{g}}$ along the desired path $\mathbf{r}$ and generates a time-varying equilibrium $\mathbf{x}^* = \boldsymbol{\ell}(\mathbf{g}, \bar{\mathbf{g}})$ for the closed-loop Hamiltonian system; (b) A local projected goal $\bar{\mathbf{g}}$ (purple dot) is generated as the furthest intersection between the local safe set $\mathcal{LS}(\mathbf{x}, \mathbf{g})$ (yellow sphere) and the path $\mathbf{r}$ (blue curve). Given $\mathbf{g}$ and $\bar{\mathbf{g}}$, a desired equilibrium $\mathbf{x}^* = \boldsymbol{\ell}(\mathbf{g}, \bar{\mathbf{g}})$ is generated for the Hamiltonian system with orientation indicated by the red, green, and blue arrows, respectively.

where $\mathbf{p}^* = \mathbf{g}$ and $\mathbf{r}_1^*, \mathbf{r}_2^*, \mathbf{r}_3^*$ are the rows of the matrix:

$$\mathbf{R}^*(\mathbf{g}, \bar{\mathbf{g}}) = \begin{cases} \mathbf{I} & \text{if } \mathbf{e}_3 \times \mathbf{c}_1 = 0 \\ [\mathbf{c}_1 \ \ \mathbf{c}_2 \ \ \mathbf{c}_3] & \text{otherwise,} \end{cases} \tag{16}$$

with $\mathbf{e}_3 = [0, 0, 1]^\top$, $\mathbf{c}_1 = (\bar{\mathbf{g}} - \mathbf{g})/\|\bar{\mathbf{g}} - \mathbf{g}\|$, $\mathbf{c}_2 = (\mathbf{e}_3 \times \mathbf{c}_1)/\|\mathbf{e}_3 \times \mathbf{c}_1\|$ and $\mathbf{c}_3 = (\mathbf{c}_1 \times \mathbf{c}_2)/\|\mathbf{c}_1 \times \mathbf{c}_2\|$. If $\mathbf{g} = \bar{\mathbf{g}}$, the most recent backup of $\mathbf{R}^*$ or $\mathbf{I}$ may be used.

Our safe tracking control design is visualized in Fig. 2. It consists of two parts: 1) a first-order reference governor system with state $\mathbf{g}$ adaptively following the local projected goal $\bar{\mathbf{g}}$ along the path $\mathbf{r}$ and 2) a closed-loop Hamiltonian system tracking the reference signal $\mathbf{x}^* = \boldsymbol{\ell}(\mathbf{g}, \bar{\mathbf{g}})$. Our main result is summarized in the following theorem.

Theorem 4 *Given a reference path $\mathbf{r}$, consider the closed-loop Hamiltonian system in (11) and the closed-loop reference governor system, $\dot{\mathbf{g}} = -k_g (\mathbf{g} - \bar{\mathbf{g}})$, with local projected goal $\bar{\mathbf{g}}$ provided in Definition 3. Suppose that the initial state $(\mathbf{x}_0, \mathbf{g}_0)$ satisfies:*

$$\Delta E\left(\mathbf{x}_0, \boldsymbol{\ell}(\mathbf{g}_0, \bar{\mathbf{g}}_0)\right) > 0, \quad \mathbf{g}_0 = \mathbf{r}(0) = \mathbf{p}(t_0) \in int(\mathcal{F}), \tag{17}$$

where $\Delta E(\mathbf{x}, \mathbf{x}^*)$ is a dynamic safety margin defined in (12). Then, the joint system state $(\mathbf{x}, \mathbf{g})$ converges to $(\boldsymbol{\ell}(\mathbf{r}(1), \mathbf{r}(1)), \mathbf{r}(1))$ without violating the constraints, i.e., $\mathbf{p}(t) \in \mathcal{F}, \forall t \geq t_0$.

Proof (sketch) Let $\Delta E(t) := \Delta E\left(\mathbf{x}(t), \boldsymbol{\ell}(\mathbf{g}(t), \bar{\mathbf{g}}(t))\right)$. Whenever $\Delta E(t) > 0$, the reference governor system proceeds along the reference path and updates the regulation point $\mathbf{x}^* = \boldsymbol{\ell}(\mathbf{g}, \bar{\mathbf{g}})$. Then, controller $\pi(\mathbf{x}, \mathbf{x}^*)$ regulates the closed-loop system (11) towards $\mathbf{x}^*$. During this process, the convergence speed is determined adaptively by the dynamic safety margin $\Delta E(t)$, fast when it is large and slow when small. By continuity, $\Delta E(t)$ cannot become negative without crossing 0 from above, at which point the governor system stops. Hence, by Prop. 2, safety is always enforced. This process continues until the augmented system is stabilized at $(\boldsymbol{\ell}(\mathbf{r}(1), \mathbf{r}(1)), \mathbf{r}(1))$ where $\bar{\mathbf{g}}$ stops changing. The complete proof can be found in the extended version (Li et al., 2021). ∎
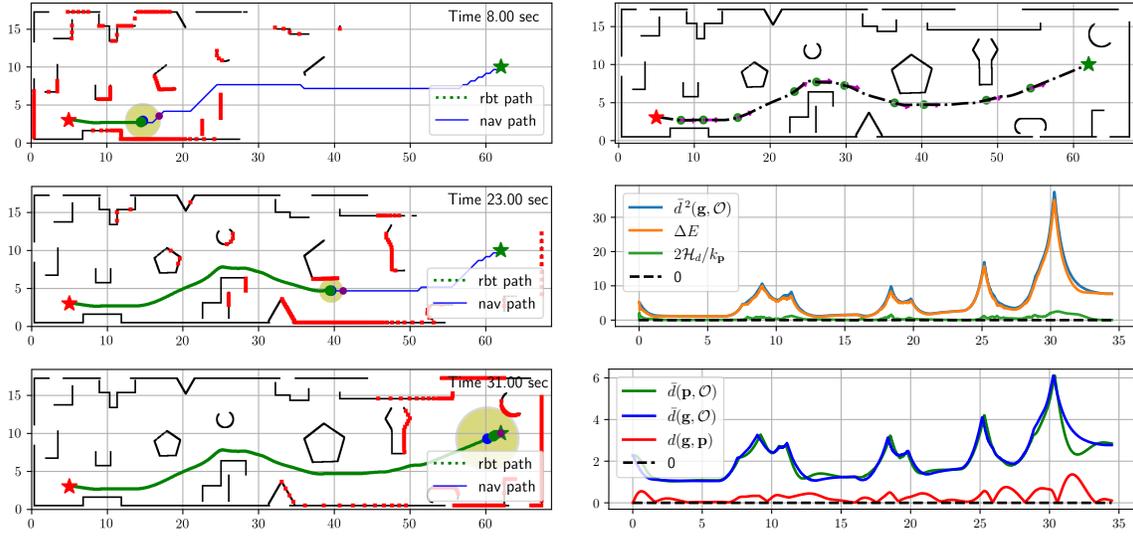
Figure 3: Safe tracking control for a hexarotor aerial robot in an unknown cluttered environment. The hexarotor (green dot) navigates from a start location (red star) to a goal location (green star) while avoiding obstacles. The obstacles are sensed by a simulated LiDAR sensor (red points). On the left are snapshots of the environment showing the robot (green dot), lidar scans (red dots), and the unsafe set $\mathcal{O}$ (black surfaces) at different times. The reference path (blue curve) is re-planned online from the governor position (blue dot) to the goal (green star) using an A^* algorithm. The local projected goal $\bar{\mathbf{g}}$ (purple dot) is computed based on the obstacle distance (gray ball) and the local safe set (yellow ball). On the upper right, the robot's heading (purple arrow) are plotted along the path. The middle right figure shows the dynamic safety margin ΔE in (12) and the scaled desired Hamiltonian $2\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*)/k_{\mathbf{p}}$ while the lower right figure plots the distance to the obstacles $d(\mathbf{p}(t), \mathcal{O})$ (green curve), showing that the safety constraints are never violated.

5. Evaluation

This section evaluates our safe tracking controller on a simulated hexarotor UAV using the learned Hamiltonian dynamics in Sec. 3.2. The task is to navigate from a start position to a goal in an environment without colliding with the obstacles. To guarantee stability and safety, all control gains must be positive. The following control gains were used with the regulation controller in Sec. 4.1: $k_p = 0.25$, $k_{\mathbf{R}} = 125\mathbf{J}$, $k_v = 0.125$, $k_{\boldsymbol{\omega}} = 10\mathbf{J}$ in (8) and $k_{\mathbf{g}} = 1.0$ with the governor control in (13). We test the controller in a challenging unknown 3D environment with complex obstacles $\mathcal{O}$. A simulated LiDAR scanner provides point cloud measurements $\mathcal{P}(t) := \{\mathbf{y}_i(t)\}$ of the surface of the unsafe set $\mathcal{O}$, depending on the system pose, with a maximum sensing range of $\beta = 30$. The distance from the governor $\mathbf{g}(t)$ to the unsafe set $\mathcal{O}$ is approximated via $\bar{d}(\mathbf{g}(t); \mathcal{O}) \approx \min_{\mathbf{y} \in \mathcal{P}(t)} \|\mathbf{g}(t) - \mathbf{y}\|$. The point clouds $\mathcal{P}(t)$ are used to construct an occupancy grid map online and a reference path $\mathbf{r}$ is replanned periodically using the A* algorithm to ensure that $\mathbf{r}(\sigma) \in \text{int}(\mathcal{F})$. In this paper, we assume the learned system dynamics are accurate and simulated a noise-free environment. We leave model uncertainty, measurement noise, and external disturbances for future work.

9
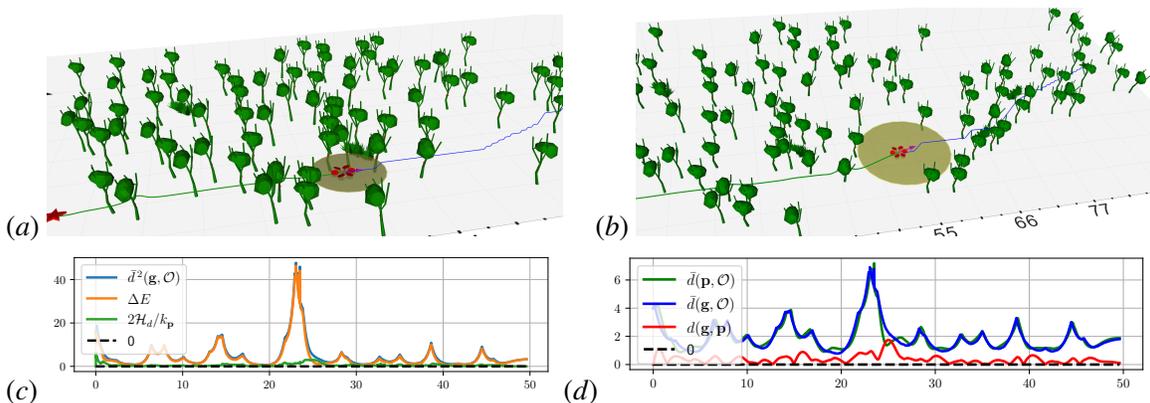
Figure 4: Safe navigation of a hexarotor robot in a forest is demonstrated in (a) and (b). The bottom plots (c) and (d) show the dynamic safety margin ΔE, the scaled desired Hamiltonian $2\mathcal{H}_d(\mathbf{x}, \mathbf{x}^*)/k_{\mathbf{p}}$, and the distance to the obstacles $d(\mathbf{p}(t), \mathcal{O})$, indicating that the safety constraints are never violated.

Fig. 3 and Fig. 4 show the behavior of the closed-loop hexarotor system in two different environments. The reference governor follows the projected goal $\bar{\mathbf{g}}$ and generates a time-varying equilibrium $\mathbf{x}^* = \ell(\mathbf{g}, \bar{\mathbf{g}})$ for the hexarotor. The dynamic safety margin $\Delta E(\mathbf{x}, \mathbf{x}^*)$ fluctuates during this process but it never becomes negative as seen in the figures. The augmented system $(\mathbf{x}, \mathbf{g})$ is controlled adaptively, slowing down when the dynamic safety margin decreases (e.g., when the robot is close to an obstacle and has large total energy $\mathcal{H}_d$) and speeding up otherwise (e.g., when the robot is far away from the obstacles or has small total energy $\mathcal{H}_d$). The simulations show that our control policy successfully drives the system from the start to the end of the reference path while avoiding sensed obstacle online, i.e., $d(\mathbf{p}, \mathcal{O})$ remains positive throughout the motion.

6. Conclusion

This paper developed a tracking controller for fully-actuated Hamiltonian systems which enables safe autonomous navigation in unknown environments. Given only a training set of system state-control trajectories, our approach estimates the system dynamics accurately using a neural ODE network and synthesizes a controller that avoids obstacles based on run-time distance measurements. Our method was demonstrated on a simulated hexarotor aerial robot navigating in complex 3D environments. Future work will focus on capturing model uncertainty and external disturbances in the design and deploying it on a hardware platform.

Acknowledgments

References

Aaron D Ames, Kevin Galloway, Koushil Sreenath, and Jessy W Grizzle. Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics. *IEEE Transactions on Automatic Control (TAC)*, 59(4), 2014a.

Aaron D Ames, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *IEEE Conference on Decision and Control (CDC)*, 2014b.

Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control (TAC)*, 62(8), 2017.

Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *IEEE European Control Conference (ECC)*, 2019.

Omur Arslan and Daniel E Koditschek. Smooth extensions of feedback motion planners via reference governors. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

Alberto Bemporad. Reference governor for constrained nonlinear systems. *IEEE Transactions on Automatic Control (TAC)*, 43(3), 1998.

Francesco Borrelli, Alberto Bemporad, and Manfred Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

José Manuel Bravo, Teodoro Alamo, and Eduardo F Camacho. Robust mpc of constrained discrete-time nonlinear systems based on approximated reachable sets. *Automatica*, 42(10), 2006.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423, 2015. doi: 10.1109/TPAMI.2013.218.

Jeffrey Delmerico and Davide Scaramuzza. A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2502–2509, 2018.

Thai Duong and Nikolay Atanasov. Adaptive control of SE(3) Hamiltonian dynamics with learned disturbance features. *arXiv preprint arXiv:2109.09974*, 2021a.

Thai Duong and Nikolay Atanasov. Hamiltonian-based neural ODE networks on the SE(3) manifold for dynamics learning and control. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021b. doi: 10.15607/RSS.2021.XVII.086.

Emanuele Garone and Marco M Nicotra. Explicit reference governor for constrained nonlinear systems. *IEEE Transactions on Automatic Control (TAC)*, 61(5), 2016.

Lars Grüne and Jürgen Pannek. *Nonlinear model predictive control*. Springer, 2017.

Sylvia L Herbert, Mo Chen, SooJean Han, Somil Bansal, Jaime F Fisac, and Claire J Tomlin. FaSTrack: A modular framework for fast and guaranteed safe motion planning. In *IEEE Conference on Decision and Control (CDC)*, 2017.

Darryl D Holm. *Geometric Mechanics*. World Scientific Publishing Company, 2008.

Juraj Kabzan, Lukas Hewing, Alexander Liniger, and Melanie N. Zeilinger. Learning-based model predictive control for autonomous racing. *IEEE Robotics and Automation Letters*, 4(4):3363–3370, 2019. doi: 10.1109/LRA.2019.2926677.

Hassan Khalil. *Nonlinear systems*. Prentice Hall, 2002.

Ilya Kolmanovsky, Emanuele Garone, and Stefano Di Cairano. Reference and command governors: A tutorial on their theory and automotive applications. In *IEEE American Control Conference (ACC)*, 2014.

Shreyas Kousik, Sean Vaskov, Fan Bu, Matthew Johnson-Roberson, and Ram Vasudevan. Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots. *The International Journal of Robotics Research (IJRR)*, 2020.

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

Zhichao Li, Omur Arslan, and Nikolay Atanasov. Fast and Safe Path-Following Control using a State-Dependent Directional Metric. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

Zhichao Li, Thai Duong, and Nikolay Atanasov. Safe autonomous navigation for systems with learned SE(3) Hamiltonian dynamics. *arXiv preprint arXiv:2112.04639*, 2021.

Anatolii Isakovich Lurie. *Analytical mechanics*. Springer Science & Business Media, 2013.

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.

Anirudha Majumdar and Russ Tedrake. Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research (IJRR)*, 36(8), 2017.

David Q Mayne, James B Rawlings, Christopher V Rao, and Pierre OM Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 2000.

Sherif AS Mohamed, Mohammad-Hashem Haghbayan, Tomi Westerlund, Jukka Heikkonen, Hannu Tenhunen, and Juha Plosila. A survey on odometry for autonomous navigation systems. *IEEE Access*, 7:97466–97486, 2019.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Multistep neural networks for data-driven discovery of nonlinear dynamical systems. *arXiv preprint arXiv:1801.01236*, 2018.

Sujit Rajappa, Markus Ryll, Heinrich H Bülthoff, and Antonio Franchi. Modeling, control and design optimization for a fully-actuated hexarotor aerial vehicle with tilted propellers. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4006–4013, 2015.

Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.

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.

Xiangru Xu. Constrained control of input–output linearizable systems using control sharing barrier functions. *Automatica*, 87:195–201, 2018.

Michael C Yip and David B Camarillo. Model-less feedback control of continuum manipulators in constrained environments. *IEEE Transactions on Robotics*, 30(4):880–889, 2014.

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