

Nonholonomic Yaw Control of an Underactuated Flying Robot With Model-Based Reinforcement Learning

Nathan O. Lambert , Craig B. Schindler , Daniel S. Drew , and Kristofer S. J. Pister 

Abstract—Nonholonomic control is a candidate to control nonlinear systems with path-dependant states. We investigate an underactuated flying micro-aerial-vehicle, the ionocraft, that requires nonholonomic control in the yaw-direction for complete attitude control. Deploying an analytical control law involves substantial engineering design and is sensitive to inaccuracy in the system model. With specific assumptions on assembly and system dynamics, we derive a Lie bracket for yaw control of the ionocraft. As a comparison to the significant engineering effort required for an analytic control law, we implement a data-driven model-based reinforcement learning yaw controller in a simulated flight task. We demonstrate that a simple model-based reinforcement learning framework can match the derived Lie bracket control – in yaw rate and chosen actions – in a few minutes of flight data, without a predefined dynamics function. This letter shows that learning-based approaches are useful as a tool for synthesis of nonlinear control laws previously only addressable through expert-based design.

Index Terms—Reinforcement learning, nonholonomic motion planning, aerial systems: mechanics and control.

I. INTRODUCTION

GENERATING control laws for novel robots is subject to high cost-per-test and imperfect dynamics models, and these deleterious effects are amplified when deriving complex, nonlinear controllers. For example, nonholonomic path planning has been applied to multiple underactuated systems, but requires precise knowledge of the controlled dynamical system and assumptions on resulting motion. In the context of novel robotics, which are often hand assembled, a perfect understanding of the dynamics is improbable. We investigate if the performance of a novel nonholonomic control law can be matched with data-driven, reinforcement learning techniques that do not require any assumptions on dynamics. Model-based reinforcement learning (MBRL) has been successfully applied to a variety of robotics


Manuscript received September 2, 2020; accepted December 12, 2020. Date of publication December 21, 2020; date of current version January 7, 2021. This letter was recommended for publication by Associate Editor S. Rathinam and Editor L. Pallottino upon evaluation of the reviewers' comments. (Corresponding author: Nathan O. Lambert.)

Nathan O. Lambert, Craig B. Schindler, and Kristofer S. J. Pister are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: nol@berkeley.edu; craig.schindler@berkeley.edu; pister@eecs.berkeley.edu).


Daniel S. Drew is with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94305-2203 USA (e-mail: ddrew73@berkeley.edu).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2020.3045930>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2020.3045930



(a) The studied robot is shown next to a U.S. quarter with the four thruster forces and the IMU orientation labelled. It masses at only 30 mg and generates thrust from a high voltage, “ion wind” [9].



(b) An illustration of Lie bracket control to produce a change in yaw. The Lie bracket is the manifold of motion between the starting and final position above, yielding $\text{yaw} = e^2[f, g](x_0)$ at $t = 4\epsilon$. Our simulations show that strict Lie bracket control can achieve angular rates of up to 13.2 degrees/s. This performance is exceeded with a model-based reinforcement learning method.

Fig. 1. Above: The novel robot we analyze, the ionocraft [9]. Below: The proposed nonholonomic yaw controller.

tasks [1]–[4]. MBRL techniques iteratively plan a sequence of actions on an data-optimized dynamics model. In this letter we ask a question: can data-driven model-based planning effectively plan multiple steps of nonholonomic control?

Nonholonomic control is most known for the example of the parallel parking problem in automobiles – defined by constraints limiting the path a system can take [5], [6]. Multiple applications of nonholonomic control in robotics demonstrate the potential to control these systems with closed loop, nonlinear control laws [7], such as the Lie bracket sequence [8]. Lie bracket control executes a specific, repeated control law in two planes

of motion to achieve motion in an only-indirectly addressable control manifold.

In this letter, we contextualize our exploration of nonholonomic control and model-based planning with the ionocraft platform, a centimeter-scale flying robot with dynamics similar to a quadrotor [9], [10] – i.e., there are four individually actuatable thrusters arranged in a quadrant. An ionocraft is an aerial vehicle which generates thrust using ionic thrusters, characterized in [11]. Quadrotors have control authority over the yaw axis via differential rotor rates that create torque about the z -axis, but the four-thruster ionocraft studied lacks this torque mechanism [12]. A Lie bracket manifold exists in the yaw direction for the studied ionocraft via repeated pitch and roll motions. The derived yaw mechanism for control via a Lie bracket approach is summarized in Fig. 1(b).

The contribution in this letter is to show the potential of data-driven, model-based RL to improve on analytical nonholonomic control in the context of a passively unstable, rapid flying robot. Specifically, in this manuscript, we 1) derive a new nonholonomic, Lie bracket control law for the ionocraft’s yaw axis and 2) demonstrate that model-based reinforcement learning improves on the analytical control policy in terms of yaw rate and crash frequency.

II. RELATED WORKS

A. Novel Aerial Vehicles

Development of novel actuators and small-scale electronics has progressed to multiple novel micro-aerial vehicles (MAVs) [13], [14]. Variants on flying robots involve robots with spinning propellers, such as a tri-rotor flying robot [15] and others with one spinning motor [16], [17]. Other robots have shown the potential for taking advantage of the flapping wing mechanism found in nature [18], [19]. Robots using ion-thrusters have only recently taken off, and are yet to be controlled in experiment [9], [10]. In contrast with automobiles or walking robots, flying robots have simpler (absence of contact forces), high-speed dynamics, but are not passively stable in the sense that they are prone to crashing, resulting in a challenging control problem.

B. Nonholonomic Robotic Control

Nonholonomic control is a sub-area of nonlinear control theory studying systems with path-dependant states [5], [6]. One method of nonholonomic control uses the Lie bracket, which involves repeated, mirrored actuation on two input manifolds [8], as detailed in Sec. III-A. Nonholonomic control is frequently applied to wheeled robots where the plane of motion must follow the direction of the wheels, leading to the parallel parking problem [20]. An extension of the 4-wheeled parallel parking problem is the controllability of two-wheeled balancing robots [21]. A novel trident snake robot uses nonholonomic control to maneuver with three arms attached to non-slide wheels [22].

There are limited studies of nonholonomic control applied to flying robots. Nonholonomic path planning is used for simulated

flying vehicles [23], but it does not consider details of low-level dynamics needed for attitude control. A space-robot with a multi-joint arm attached to a main body uses nonholonomic control with a Lyapunov function for path-planning [24] and nonholonomic redundancy to avoid joint limits [25]. Another contactless-medium of motion is swimming, and nonholonomic control is used to imitate carangiform swimming of a planar, rigid fish-robot [26] and to control a AUV [27]. A novel nonholonomic controller called “wriggle-steering” uses repeated and mirrored sinusoidal oscillations – like a Lie bracket in a frequency domain – to yaw a micro-flapping wing robot [28]. Our work extends the application of nonholonomic, Lie bracket control to that of low-level control of a flying rigid-body, with challenging, rapid dynamics.

C. Model-Based Reinforcement Learning

Model-based reinforcement learning (MBRL) has emerged as a functional candidate for robotic control in a data-efficient manner [1]–[3], [29] – this letter extends the functionality of MBRL to nonholonomic planning for flying robots. Visual MBRL (i.e., control from pixels) has been shown to work in a simple, simulated nonholonomic car task with Linear Quadratic Regulator (LQR) control and variational inference [30]. Additionally, MBRL has been shown to learn to control a micro-quadrotor with no assumption on dynamics [31] and with suspended payloads modifying dynamics [32]. MBRL for attitude control of quadrotors is encouraging for data-driven control of a novel MAV, and this work extends the results to show the capability of MBRL algorithms to plan complex multi-step action sequences for flying robots.

III. BACKGROUND

A. Lie Bracket Control

For a nonlinear, dynamical system with two inputs, $v = [v_1 \ v_2]$, a Lie bracket control law is defined by a repeated sequence of motions for a small time, ϵ , to generate motion in a new manifold of motion, $[f, g]$, called the Lie bracket vector field. Consider the following system,

$$\dot{x} = f(x)v_1 + g(x)v_2 \quad (1)$$

The Lie bracket control law is motivated by a series of Taylor series expansions along the vector fields of each input f, g , so that a third manifold of motion emerges, $[f, g]$. Specifically, consider the input sequence below,

$$v(t) = \begin{cases} [1, 0] & t \in [0, \epsilon) \\ [0, 1] & t \in [\epsilon, 2\epsilon) \\ [-1, 0] & t \in [2\epsilon, 3\epsilon) \\ [0, -1] & t \in [3\epsilon, 4\epsilon) \end{cases} \quad (2)$$

After $t = 4\epsilon$, the state can be expressed as (shown in [33])

$$x(4\epsilon) = x_0 + \epsilon^2[f, g](x_0) + \mathcal{O}(\epsilon^3), \quad (3)$$

$$\text{where } [f, g] = \frac{df}{dx}g - \frac{dg}{dx}f. \quad (4)$$

For a visualization of the Lie bracket motion in a phase plane caused by two inputs, see Fig. 1 b.

B. Model-Based Reinforcement Learning

Model-based reinforcement learning (MBRL) follows the framework of an agent interacting in an environment, learning a model of said environment, and then leveraging the model for control [34]. Specifically, the agent acts in a Markov Decision Process (MDP) governed by a transition function $x_{t+1} = f(x_t, u_t)$ and returns a reward at each step $r(x_t, u_t)$. With a collected dataset $\mathcal{D} := \{x_i, u_i, x_{i+1}, r_i\}_{i=1}^N$, the agent learns a model, $x_{t+1} = f_\theta(x_t, u_t)$ to minimize the negative log-likelihood of the transitions. We employ sample-based model-predictive control (MPC) using the learned dynamics model, which optimizes the expected reward over a finite, recursively predicted horizon, τ , from a set of actions sampled from a uniform distribution $\mathcal{U}(a)$, as

$$u^* = \arg \max_{u_{t:t+\tau}} \sum_{i=0}^{\tau} r(\hat{x}_{t+i}, u_{t+i}), \quad (5)$$

$$s.t. \quad \hat{x}_{t+1} = f_\theta(\hat{x}_t, u_t). \quad (6)$$

IV. IONOCRAFT MODEL

In this letter, we model and study an ionocraft [9], [10]; a flying robot with four individually-addressable electrohydrodynamic thrusters, as shown in Fig. 1. The robot is approximately 4cm² and masses 30 mg not including the sensor payload. Electrohydrodynamic thrust, summarized in [11], is produced by the momentum-transferring collisions between ions drifting in an applied electric field and the particles in a neutral fluid. The forces, $F_k(i_k, \beta_k)$, are models for the electrohydrodynamic force from each thruster [9], as

$$F_{ion,k} = \beta_k \frac{i_k d}{\mu}. \quad (7)$$

Here, i is the ion current, $d = 500\mu\text{m}$ is the air-gap and $\mu = 2 \cdot 10^{-4} \text{m}^2/\text{V}$ is the ion mobility of N_2^+ in dry air. Fitting a β_k , which represents a deviation in force from that expected by an idealized one-dimensional derivation, to each thruster allows more flexibility in model fitting and simulation. Past work shows $\beta \in [0.5, 0.7]$ for the region of currents $i \in [0, .5] \text{mA}$ corresponding to voltages $v \in [1650, 2100] \text{V}$ [9].

The ionocraft dynamics model follows a 12 state rigid-body dynamics with a state $x = [X \ Y \ Z \ \psi \ \theta \ \phi \ v_x \ v_y \ v_z \ \omega_x \ \omega_y \ \omega_z]$, and input $u = [F_1 \ F_2 \ F_3 \ F_4]$ mN, centered around an unstable equilibrium point $x^* = \mathbf{0}$, which are simulated at a dynamics frequency of 1kHz, a control frequency of 100 Hz to match the experimental setup used in [35], and Gaussian state noise in each dimension, $\sigma \sim \mathcal{N}(0, 0.01)$:

$$\begin{bmatrix} \dot{X} \\ \dot{Y} \\ \dot{Z} \end{bmatrix} = Q_{B/I} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix}, \quad (8)$$

$$\begin{bmatrix} \dot{\psi} \\ \dot{\theta} \\ \dot{\phi} \end{bmatrix} = W^{-1} \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix}, \quad (9)$$

$$\begin{bmatrix} \dot{v}_x \\ \dot{v}_y \\ \dot{v}_z \end{bmatrix} = \frac{1}{m} \begin{bmatrix} F_x \\ F_y \\ F_z \end{bmatrix} - \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix}, \quad (10)$$

$$I_B \begin{bmatrix} \dot{\omega}_x \\ \dot{\omega}_y \\ \dot{\omega}_z \end{bmatrix} = \begin{bmatrix} \tau_x \\ \tau_y \\ \tau_z \end{bmatrix} - \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix} I_B \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix} \quad (11)$$

The inverse Wronskian $W^{-1} \in \mathbb{R}^{3 \times 3}$ and the body-to-inertial transformation $Q_{B/I} \in \mathbb{R}^{3 \times 3}$ defined with the trigonometric functions abbreviated as $c := \cos$, $s := \sin$, and $t := \tan$:

$$Q_{B/I} = \begin{bmatrix} c(\theta)c(\psi) & c(\psi)s(\theta)s(\phi) & s(\phi)s(\psi) \\ & -c(\phi)s(\psi) & +c(\phi)c(\psi)s(\theta) \\ c(\theta)s(\psi) & c(\phi)c(\psi) & c(\phi)s(\theta)s(\psi) \\ & +s(\theta)s(\phi)s(\psi) & -c(\psi)s(\phi) \\ -s(\theta) & c(\theta)s(\phi) & c(\theta)c(\phi) \end{bmatrix} \quad (12)$$

$$W^{-1} = \frac{1}{c(\theta)} \begin{bmatrix} 0 & s(\phi) & c(\phi) \\ 0 & c(\phi)c(\theta) & -s(\phi)c(\theta) \\ c(\theta) & s(\phi)s(\theta) & c(\phi)s(\theta) \end{bmatrix}. \quad (13)$$

The simulator emulates low-level control by actuating the ionocraft with voltages, which map to currents and forces via Eq. (7). Then, the rigid-body model incorporates torques and forces as an ideal point-force from the center of the thruster, $l = 1$ cm, from the center-of-mass.

$$\begin{bmatrix} F_z \\ \tau_z \\ \tau_y \\ \tau_x \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ -l & -l & l & l \\ -l & l & l & -l \end{bmatrix} \begin{bmatrix} F(i_1, \beta) \\ F(i_2, \beta) \\ F(i_3, \beta) \\ F(i_4, \beta) \end{bmatrix}. \quad (14)$$

The important observation for yaw control is that the z -axis torque is totally decoupled from the individual thrusters. Neglecting the damping forces in low velocity maneuvers such as hovering, a corresponding subset of the Jacobian linearization of the ionocraft approximated around equilibrium, $x^* = \mathbf{0}$, is

$$\dot{\psi} = \omega_z, \quad \dot{\theta} = \omega_y, \quad \dot{\phi} = \omega_x + \omega_z, \quad (15)$$

$$\dot{\omega}_x = \frac{\tau_x}{I_{xx}}, \quad \dot{\omega}_y = \frac{\tau_y}{I_{yy}}, \quad \dot{\omega}_z = \frac{\tau_z}{I_{zz}}. \quad (16)$$

A direct way to see a loss of yaw-actuation is with a linearized controllability matrix, $\tilde{W}_c(\tilde{A}, \tilde{B})$,

$$\tilde{W}_c = [\tilde{B} \ \tilde{A}\tilde{B} \ \dots \ \tilde{A}^{n-1}\tilde{B}]. \quad (17)$$

A robot can actuate all directions if this matrix is full rank. The matrix \tilde{W}_c is of rank $n - 1$, with a corresponding uncontrollable mode of ω_z , but this also propagates to ψ . This loss of controllability formally shows the need for a nonlinear yaw control law.

V. NONHOLONOMIC YAW CONTROL

In this section we derive the analytical Lie bracket controller and compare it to a data-driven MBRL approach for yaw control. With both methods, the goal of the approach is to maximize the yaw, ψ , without causing substantial flight disturbance (e.g. reaching the stop condition in the simulator of the magnitude of pitch, θ , or roll, ϕ , above 45 degrees). Code, video, and an expanded manuscript are available on the website: <https://sites.google.com/berkeley.edu/mbrl-ionocraft/>.

A. Analytical Control: Lie Bracket

Here we map the Lie bracket formulation outlined in Sec. III-A to an open-loop yaw controller for the ionocraft. First, consider the Euler rates defined in Eq. (11), with the trigonometric functions abbreviated.

$$\dot{\psi} = \frac{s(\phi)}{c(\theta)}\omega_y + \frac{c(\phi)}{c(\theta)}\omega_z, \quad (18)$$

$$\dot{\theta} = c(\phi)\omega_y - s(\phi)\omega_z, \quad (19)$$

$$\dot{\phi} = \omega_x + s(\phi)t(\theta)\omega_y + c(\phi)\omega_z. \quad (20)$$

To study as nonlinear control, expand the differential equations for Euler rates with angular rates as inputs. In order to assume the angular rates are inputs, we must show that they are directly and independently actuated. To do so, consider the angular rate differential equations shown below, assuming the inertial matrix is diagonal: $I_B = \text{diag}(I_{xx}, I_{yy}, I_{zz})$.

$$\dot{\omega}_x = \frac{1}{I_{xx}} (\tau_x + \omega_z\omega_y I_{yy} - \omega_y\omega_z I_{zz}), \quad (21)$$

$$\dot{\omega}_y = \frac{1}{I_{yy}} (\tau_y + \omega_z\omega_x I_{xx} - \omega_x\omega_z I_{zz}), \quad (22)$$

$$\dot{\omega}_z = \frac{1}{I_{zz}} (\tau_z + \omega_y\omega_x I_{xx} - \omega_x\omega_y I_{yy}). \quad (23)$$

The following assumptions are needed to prove the existence of the Lie bracket yaw controller:

- 1) For a symmetric body in the XY -plane, $I_{xx} = I_{yy}$, so the cross terms of $\dot{\omega}_z$ are equal, $\omega_y\omega_x I_{xx} = \omega_x\omega_y I_{yy}$.
- 2) The robot starts from rest, $\omega_z(t_0) = 0$, so the cross terms in $\dot{\omega}_x, \dot{\omega}_y$ are zero initially.
- 3) With the lack of yaw coupling in thrusters, the robot applies 0 torque about the z axis, $\tau_z = 0$.

Combining assertions 2) and 3) enforces $\omega_z = 0 \forall t$. With 1), this means that $\dot{\omega}_x, \dot{\omega}_y$ are fully determined by τ_x, τ_y , so we can model the flight dynamics as a nonlinear system with two manifolds of action, as in Eq. (1).

Therefore, for the flying robot define the state space with the following state variables and vector fields:

$$\mathbf{x} = [\psi \quad \theta \quad \phi], \quad \tilde{\mathbf{v}} = [\omega_x \quad \omega_y \quad \omega_z], \quad (24)$$

$$\begin{bmatrix} \dot{\psi} \\ \dot{\theta} \\ \dot{\phi} \end{bmatrix} = f(x)\omega_x + g(x)\omega_y + h(x)\omega_z. \quad (25)$$

TABLE I

DIFFERENT ASSEMBLY CONFIGURATIONS AND PARAMETERS SHOWING THE CHANGE OF INERTIAL PROPERTIES FROM FABRICATION ERROR IN IMU PLACEMENT – BREAKING THE SYMMETRIC BODY ASSUMPTIONS NEEDED TO SAFELY APPLY LIE BRACKET YAW CONTROL. THE INERTIAL VALUES ARE REPORTED IN G/MM^2 AND THE CROSS TERMS ARE SMALL, I.E., $\{I_{XY}, I_{XZ}, \dots\} < .01$

No IMU		IMU Center		IMU 5mm x-error	
Mass	26mg	Mass	46mg	Mass	46mg
I_{XX}	1.967	I_{XX}	1.984	I_{XX}	2.262
I_{YY}	1.967	I_{YY}	1.983	I_{YY}	1.983
I_{ZZ}	3.775	I_{ZZ}	3.804	I_{ZZ}	4.083

Define the vector fields from Eq. (18), Eq. (19), and Eq. (20).

$$f(x) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} g(x) = \begin{bmatrix} \frac{s(\phi)}{c(\theta)} \\ c(\phi) \\ s(\phi)t(\theta) \end{bmatrix} h(x) = \begin{bmatrix} \frac{c(\phi)}{c(\theta)} \\ -s(\phi) \\ c(\phi) \end{bmatrix}. \quad (26)$$

Making the assumptions above guarantees that $h(x)$ does not affect the differential equation, and there is no direct yaw actuation. Consider two Lie brackets $[f, g]$ and $[g, f]$, as defined in Eq. (4).

Computing the Jacobian of our dynamics f and g yields,

$$\frac{df}{dx} = 0_{3 \times 3}, \quad \frac{dg}{dx} = \begin{bmatrix} 0 & \frac{2s(\phi)s(\theta)}{c(2\theta)+1} & \frac{c(\phi)}{c(\theta)} \\ 0 & 0 & -s(\phi) \\ 0 & s(\phi)c^2(\theta) & c(\phi)t(\theta) \end{bmatrix}. \quad (27)$$

Finally, the Lie bracket is computed as:

$$[f, g](x) = -\frac{dg}{dx}f(x) = -\begin{bmatrix} -\frac{c(\phi)}{c(\theta)} \\ s(\phi) \\ -c(\phi)t(\theta) \end{bmatrix}. \quad (28)$$

The final assumption is that the control law will be applied near hover, $\theta, \phi \approx 0$. This approximation shows controllability in the yaw-direction with low pitch or roll drift captured in the higher order terms, $\mathcal{O}(\epsilon^3)$ in Eq. (4) – completing attitude control for the ionocraft.

$$[f, g](x) \Big|_{\theta=0, \phi=0} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad (29)$$

Intuitively, the robot actuates yaw by repeated, mirrored pitch-roll oscillations, but the assumptions on symmetry are highly sensitive to the reality of imperfect assembly outlined in Table I.

Translating the input vector, v , used in the Lie bracket derivation to an input of motor forces, u , is done by setting and holding specific angular accelerations. Given a desired angular rate, the controller can set the angular rate with a simple force-to-torque model. For example, the angular acceleration about the x -axis follows: $\dot{\omega}_x = \frac{l}{I_{xx}} \cdot (F_1 + F_2 - F_3 - F_4)$, where $l = 1$ cm is the length from the center of mass to the center of the thruster. The control setup and simulator use specific application of motor forces rather than setting an desired angular rate, so we define the actions used by the Lie bracket sequence to yaw the robot in Eq. (30). Applying a inner-loop controller to set the angular

TABLE II

THE DIFFERENT YAW RATES ACHIEVED BY A LIE BRACKET CONTROLLER WITH DIFFERENT VALUES OF THE TIME-HOLD, ϵ , FOR EACH ACTION IN THE SEQUENCE. WITHOUT ANY FEEDBACK CONTROL ON PITCH OR ROLL, THE LIE BRACKET CONTROLLER DIVERGES AND REACHES THE STOP CONDITION OF PITCH OR ROLL GREATER THAN 45 DEGREES. THERE IS A QUADRATIC INCREASE IN YAW-RATE CORRESPONDING TO THE TIME-LENGTH EACH ACTION IN THE SEQUENCE (FOLLOWING EQ. (4))
 $\mathbf{u}_{\text{pitch}}^+$, $\mathbf{u}_{\text{roll}}^+$, $\mathbf{u}_{\text{pitch}}^-$, $\mathbf{u}_{\text{roll}}^-$, AT A COST OF FASTER DIVERGENCE OF ATTITUDE STABILITY

Lie Bracket ϵ (s)	Yaw Rate (deg/s)	Simulation Stop (s)
0.01	2.0	7.52
0.02	3.9	3.78
0.03	5.9	2.51
0.04	7.9	1.88
0.06	11.5	1.30
0.08	13.2	0.94

rates would aid the Lie bracket controller and aid with fabrication variance, but the current design of the robot does not have this capability. Consider \mathbf{u}_y^x , where x corresponds to a direction $\in \{+, -\}$ of the dimension $y \in \{\text{pitch}, \text{roll}, \text{equil.}\}$.

$$\begin{aligned}
 \mathbf{u}_{\text{pitch}}^+ &= [0.15, 0.05, 0.05, 0.15] \text{ mN} \\
 \mathbf{u}_{\text{roll}}^+ &= [0.15, 0.15, 0.05, 0.05] \text{ mN} \\
 \mathbf{u}_{\text{pitch}}^- &= [0.05, 0.15, 0.15, 0.05] \text{ mN} \\
 \mathbf{u}_{\text{roll}}^- &= [0.05, 0.05, 0.15, 0.15] \text{ nM} \\
 \mathbf{u}_{\text{equil}} &= [0.1, 0.1, 0.1, 0.1] \text{ mN}
 \end{aligned} \tag{30}$$

These actions are designed given flight mass of 50mg to assume a payload such as an IMU and a thrust coefficient, $\beta = 0.6$ (equilibrium is the action with all motors set to output of $F_i = \frac{mg}{4}$). A hard-coded sequence of these actions at 100 Hz ($\epsilon = 0.01$) achieves a baseline yaw rate of 1 degrees/sec. This yaw rate can be increased by a corresponding increase in the time that each action is held, ϵ . Table II shows that the raw rate can increase with the Lie bracket controller up to 13.2 degrees/s, but faster yaw actuation causes faster instability in pitch and roll. To slow divergence and get peak yaw-rate with the Lie bracket, we remove all state noise, $\sigma = 0$, and zero the initial state, $\mathbf{x}(t_0) = 0$.

Translating control methods to robotic hardware strongly correlates with the practical value of a method for control synthesis. Lie bracket control is based on careful assumptions of symmetry in movement and near-equilibrium flight, which can easily break down due to environmental disturbances or nonideality in robot assembly. Specifically, the ionocraft needs an on-board MPU-9250 inertial measurement device (IMU) placed in the center for state-space measurement. This assembly is done by hand and subject to errors, which can have substantial alteration on the inertial properties of the robot – summarized in Table I. The hand-assembly also conveys an increased cost-per-test, so the ideal method for generating a controller would be able to generalize across inaccurate dynamics and be sample efficient.

B. Data-driven Control: MBRL

The main shortcoming of analytical methods is the strong assumptions on system dynamics, so we explore model-based

reinforcement learning. While free from assumptions on system dynamics, the MBRL framework has a new cost of environmental interactions that we quantify and consider. The predictive dynamics model is a simple feed-forward neural network with two hidden layers, trained with the Adam optimizer for 17 epochs with a learning rate of 0.0025 and a batch size of 18. A parameter that can strongly affect performance is the prediction horizon of the MPC, and in this letter, we use a horizon of $\tau = 5$. The number of evaluated actions is $N = 500$ and the actions for each thruster are sampled from uniform force distributions $F_i \in \mathcal{U}(0, 0.3)\text{mN}$, as

$$\mathbf{u}_i = [F_1 \ F_2 \ F_3 \ F_4] \text{ mN}. \tag{31}$$

In this section we detail the design decisions that must be made to achieve nonholonomic behavior with a model-based controller, which primarily entails generating a suitable reward function.

A simple formulation for achieving yaw in a reinforcement learning framework would be to encourage maximizing the magnitude of yaw, such as $r(\mathbf{x}, \mathbf{u}) = \psi$, but this reward function does not penalize a loss of attitude stability. In practice, this reward function quickly diverges and hits the simulator pitch, roll stop condition of $|\theta| > 45$ degrees or $|\phi| > 45$ degrees. To mitigate this, the objective becomes a dual optimization with the reward function balancing the reward of yaw versus the cost of pitch and roll. Simply appending a scaled squared cost on pitch and roll (e.g. $r(\mathbf{x}, \mathbf{u}) = \psi^2 - \lambda(\theta^2 + \phi^2)$) does not achieve stable yaw control because as yaw increases globally, the controller accepts a corresponding increase in pitch and roll to increase reward – eventually the robot reaches high angles and ends the simulation early in 30% of trials, even when trained on minutes of flight data.

To alleviate this, sliding mode control [36] can ensure attitude stability while maximizing yaw rate when in a stable region of the state space, as

$$r(\mathbf{x}, \mathbf{u}) = \begin{cases} |\psi| & (|\phi| < \eta) \cap (|\theta| < \eta), \\ -(\theta^2 + \phi^2) & \text{else.} \end{cases} \tag{32}$$

This introduces a tune-able hyperparameter, the attitude window, η , weighting the balance between attitude stability and yaw rate that was set to 10 degrees during our experiments. With this reward function, agents never reach the simulation stop condition when trained on at least 20 s of flight data.

In order to test the robustness of the MBRL approach to manufacturing and environmental disturbances that substantially degrade the performance of the Lie bracket control law, we randomized the inertial properties to $\pm 15\%$ the nominal value and sample initial Euler angles $\phi, \theta \in [-22.5, 22.5]$ degrees. The performance with this randomization matches closely to that with initial state set to zero, $\mathbf{x}_0 = 0$, and symmetric dynamics, showing the ability for the MBRL approach to adapt to variable environments, shown in Fig. 3. Each trial represents a maximum of 1000 training steps, so the full training set after 10 trials represents just 100 s of training data. The learning agents learn (with and without assembly and environmental variation) mean behavior just below the Lie bracket baseline (in variation free environment), and the MBRL agents have many trials with a substantially greater yaw rate. A simulated flight, where a MBRL

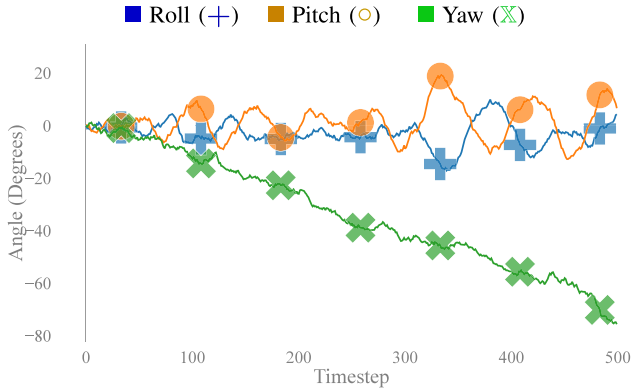


Fig. 2. The Euler angle response with MBRL-MPC of the ionocraft when maximizing the yaw rate. The out of phase pitch and roll oscillations mirror that of the Lie bracket control, where directly actuating pitch and roll is needed to generate movement in the Yaw plane.

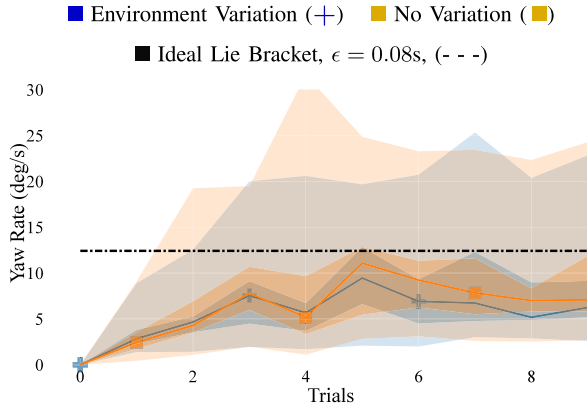


Fig. 3. The median reward ($N_{\text{robot}} = 25$), with 65th, 95th percentiles shaded, per trial when learning yaw control for the simulated ionocraft with a 15% variation on inertial properties and a wide range of initial states. The agent with environmental variation closely matches the one with symmetric dynamics and stable initial conditions, highlighting the ability for MBRL to cope with realistic hardware variation. One trial is equal to 1000 environment steps. The agent learns to stably actuate yaw, with many trials with a yaw rate over the peak derived Lie bracket at 13.2 degrees/s.

agent uses pitch and roll oscillations similar to a Lie bracket control law is shown in Fig. 2.

C. Learning to Mimic the Lie bracket

An important point when considering the learned behavior is if it mimics the Lie bracket control, or does the MBRL algorithm develop unintended, simulation-exploiting maneuvers. To test this, we train a *more constrained version* of the MBRL algorithm with a constrained action space of the actions used in Lie bracket control and equilibrium:

$$u \sim [\mathbf{u}_{\text{pitch}}^+, \mathbf{u}_{\text{roll}}^+, \mathbf{u}_{\text{pitch}}^-, \mathbf{u}_{\text{roll}}^-, \mathbf{u}_{\text{equil}}]. \quad (33)$$

The standard Lie bracket formulation actuates each action for one time step to mitigate divergence of other states from higher order interactions. In this experiment, the MBRL algorithm closely resembles the action selection of a repeated Lie bracket with some sampling noise – recall that the controller only runs with a planning horizon of $\tau = 5$. We compare the actions selected by the MBRL algorithm and a Lie bracket sequence

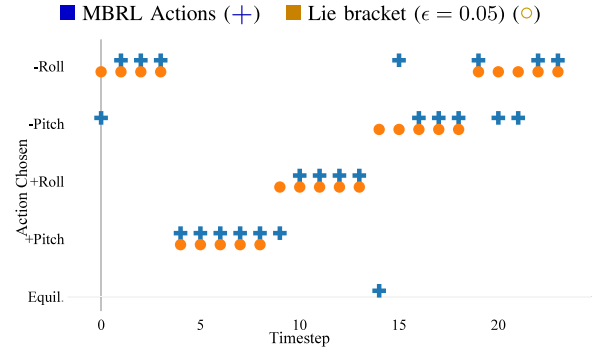


Fig. 4. A comparison of the chosen actions when the learning-based approach can only choose from the actions used in the Lie bracket from Sec. V-A and an equilibrium action. The actions represent the first 25 actions of a trajectory from a nonzero initial pitch and roll with symmetric inertial matrices using the sampling-based MPC.

repeating each action 5 times ($\epsilon = 0.05$) in Fig. 4, showing the ability for a data-driven approach to closely mimic a carefully derived analytical controller.

VI. DISCUSSION & FUTURE WORK

A model-based model-predictive controller learned to achieve a maximum yaw rate of up to 31 degrees/s by closely mimicking the Lie bracket control law. It is important to note that the MBRL controller achieves this yaw rate in a full trial, and does not cause diverging attitude stability like that of the high- ϵ Lie bracket controllers. With assembly variation and random initial states, the Lie bracket controller reaches the simulator stop condition (pitch or roll above 45degrees represents an unrecoverable state) in over 90% of trials ($N = 100$) before reaching an episode length of 1000 points (10 s) – a secondary controller preventing divergence is needed for practical application. The MBRL MPC can maintain stability and yaw actuation in the presence of assembly and environmental uncertainty, and would not need other controllers to integrate future tasks.

The means by which controllers are deployed is important to the practicality of novel methods, providing additional points of comparison for the use of the Lie-bracket or MBRL controllers with a real ionocraft. The Lie-bracket controller would be one level in a hierarchical feedback controller, requiring more expert design, and the MBRL algorithm would need a modification of the cost function to include time-varying trajectory goals. Translating the MBRL MPC to use onboard computation is an important avenue of future work, which is being explored by using advances in imitation and supervised learning to distill the MPC to a neural network policy – which should be capable of running well above 100 Hz without a GPU. The ionocraft’s movement in the yaw axis is currently constrained by the power tethers required for flight, delaying application of our method to the experimental platform [9]. Other sources of variation, such as different forces per thruster, as modeled by β_k in Eq. (7), airframe flexibility, and robot-specific drag models, are not studied and could further incentivize data-driven techniques for this robot. Integrating additional sources of uncertainty into the simulator would improve fidelity of both the Lie bracket and the MBRL.

VII. CONCLUSION

This letter provided a case study in synthesising a yaw controller for a novel flying robot, the ionocraft. Generating complex, nonlinear controllers on costly robots is a trade-off between engineer time and experimental cost. Lie brackets, and other closed-form nonholonomic controllers, require non-trivial dynamics model understanding, which is a level of expertise often lacking when deploying robotic systems. In this letter we characterized how a proven MBRL approach can closely match a derived controller at the low-level of action selection and in performance. While translating the nonholonomic control law or model-based RL to real experiments poses high risk for different reasons – namely a trade off between expert time and costly environmental interactions – this letter shows that an encouraging new path exists in data-driven control for roboticists that is not dependant on familiarity with a new platform.

ACKNOWLEDGMENT

The authors would like to thank Claire Tomlin, Francesco Borrelli, and Roberto Calandra for helpful discussions and the Berkeley Sensor & Actuator Center for their support. Thank you to Ryan Koh for helping with related code.

REFERENCES

- [1] M. Deisenroth and C. E. Rasmussen, “Pilco: A model-based and data-efficient approach to policy search,” in *Proc. 28th Int. Conf. Mach. Learn.*, 2011, pp. 465–472.
- [2] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep reinforcement learning in a handful of trials using probabilistic dynamics models,” in *Proc. Neural Inf. Process. Syst.*, 2018, pp. 4754–4765.
- [3] G. Williams *et al.*, “Information theoretic MPC for model-based reinforcement learning,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 1714–1721.
- [4] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, “Deep dynamics models for learning dexterous manipulation,” in *Proc. Conf. Robot. Learn., Ser. Mach. Learn. Res.*, L. P. Kaelbling, D. Kragic, and K. Sugiura, Eds., vol. 100, Oct.–Nov. 2020, pp. 1101–1112.
- [5] I. Kolmanovsky and N. H. McClamroch, “Developments in nonholonomic control problems,” *IEEE Control Syst. Mag.*, vol. 15, no. 6, pp. 20–36, Dec. 1995.
- [6] J.-J. E. Slotine *et al.*, *Applied Nonlinear Control*. Englewood Cliffs, NJ, USA: Prentice Hall, 1991, vol. 199, no. 1.
- [7] S. D. Kelly and R. M. Murray, “Geometric phases and robotic locomotion,” *J. Robot. Syst.*, vol. 12, no. 6, pp. 417–431, 1995.
- [8] F. C. Park, J. E. Bobrow, and S. R. Ploen, “A lie group formulation of robot dynamics,” *Int. J. Robot. Res.*, vol. 14, no. 6, pp. 609–618, 1995.
- [9] D. S. Drew, N. O. Lambert, C. B. Schindler, and K. S. Pister, “Toward controlled flight of the ionocraft: A flying microrobot using electrohydrodynamic thrust with onboard sensing and no moving parts,” *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 2807–2813, 2018.
- [10] H. K. Hari Prasad, R. S. Vaddi, Y. M. Chukewad, E. Dedic, I. Novosselov, and S. B. Fuller, “A laser-microfabricated electrohydrodynamic thruster for centimeter-scale aerial robots,” *PLoS One*, vol. 15, no. 4, p. e0231362, 2020.
- [11] K. Masuyama and S. R. Barrett, “On the performance of electrohydrodynamic propulsion,” in *Proc. Roy. Soc. Math., Phys. Eng. Sci.*, vol. 469, no. 2154, 2013, Art. no. 20120623.
- [12] R. Mahony, V. Kumar, and P. Corke, “Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor,” *IEEE Robot. Automat. Mag.*, vol. 19, no. 3, pp. 20–32, Sep. 2012.
- [13] G. Cai, J. Dias, and L. Seneviratne, “A survey of small-scale unmanned aerial vehicles: Recent advances and future development trends,” *Unmanned Syst.*, vol. 2, no. 2, pp. 175–199, 2014.
- [14] D. S. Drew, B. Kilberg, and K. S. Pister, “Future mesh-networked pico air vehicles,” in *Proc. IEEE Int. Conf. Unmanned Aircr. Syst.*, 2017, pp. 1075–1082.
- [15] J.-T. Zou, K.-L. Su, and H. Tso, “The modeling and implementation of tri-rotor flying robot,” *Artif. Life Robot.*, vol. 17, no. 1, pp. 86–91, 2012.
- [16] M. Piccolino and M. Yim, “Piccolissimo: The smallest micro aerial vehicle,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 3328–3333.
- [17] W. Zhang, M. W. Mueller, and R. D’Andrea, “Design, modeling and control of a flying vehicle with a single moving part that can be positioned anywhere in space,” *Mechatronics*, vol. 61, pp. 117–130, 2019.
- [18] R. Wood, R. Nagpal, and G.-Y. Wei, “Flight of the robobees,” *Sci. Amer.*, vol. 308, no. 3, pp. 60–65, 2013.
- [19] G. De Croon, M. Perçin, B. Remes, R. Ruijsink, and C. De Wagter, “The delfly,” *Dordrecht: Springer Netherlands.*, vol. 10, pp. 978–94, 2016.
- [20] I. E. Paromtchik and C. Laugier, “Autonomous parallel parking of a nonholonomic vehicle,” in *Proc. IEEE Conf. Int. Vehicles*, 1996, pp. 13–18.
- [21] A. Salerno and J. Angeles, “On the nonlinear controllability of a quasi-holonomic mobile robot,” in *Proc. IEEE Int. Conf. Robot. Automat. (Cat. No. 03CH37422)*, vol. 3, 2003, pp. 3379–3384.
- [22] M. Ishikawa, “Trident snake robot: Locomotion analysis and control,” *IFAC Proc. Volumes*, vol. 37, no. 13, pp. 895–900, 2004.
- [23] E. Masehian and N. Mohamadnejad, “Path planning of nonholonomic flying robots using a new virtual obstacle method,” in *Proc. IEEE 3rd RSI Int. Conf. Robot. Mechatronics*, 2015, pp. 612–617.
- [24] Y. Nakamura and R. Mukherjee, “Nonholonomic path planning of space robots via bi-directional approach,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 1990, pp. 1764–1769.
- [25] Y. Nakamura and R. Mukherjee, “Exploiting nonholonomic redundancy of free-flying space robots,” *IEEE Trans. Robot. Automat.*, vol. 9, no. 4, pp. 499–506, 1993.
- [26] K. A. Morgansen, V. Duidam, R. J. Mason, J. W. Burdick, and R. M. Murray, “Nonlinear control methods for planar carangiform robot fish locomotion,” in *Proc. IEEE Int. Conf. Robot. Automat. (Cat. No. 01CH37164)*, vol. 1, 2001, pp. 427–434.
- [27] Z. M. Zain *et al.*, “A nonholonomic control method for stabilizing an X4-AUV,” *Artif. Life Robot.*, vol. 16, no. 2, p. 202, 2011, doi: [10.1007/s10015-011-0918-8](https://doi.org/10.1007/s10015-011-0918-8).
- [28] S. B. Fuller, J. P. Whitney, and R. J. Wood, “Rotating the heading angle of underactuated flapping-wing flyers by wriggle-steering,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 1292–1299.
- [29] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to trust your model: Model-based policy optimization,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 12 498–12 509.
- [30] M. Zhang, S. Vikram, L. Smith, P. Abbeel, M. Johnson, and S. Levine, “Solar: Deep structured representations for model-based reinforcement learning,” in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 7444–7453.
- [31] N. O. Lambert, D. S. Drew, J. Yaconelli, S. Levine, R. Calandra, and K. S. Pister, “Low-level control of a quadrotor with deep model-based reinforcement learning,” *IEEE Robot. Automat. Lett.*, vol. 4, no. 4, pp. 4224–4230, 2019.
- [32] S. Belkhal, R. Li, G. Kahn, R. McAllister, R. Calandra, and S. Levine, “Model-based meta-reinforcement learning for flight with suspended payloads,” 2020, *arXiv:2004.11345*.
- [33] H. K. Khalil and J. W. Grizzle, *Nonlinear systems*. Prentice hall Upper Saddle River, NJ, vol. 3, 2002.
- [34] R. S. Sutton and A. G. Barto, *Reinforcement learn.: An Introduction*. MIT press, 2018.
- [35] D. Drew, *The Ionocraft: Flying Microrobots with no Moving Parts*. eScholarship, University of California, 2018.
- [36] C. Edwards and S. Spurgeon, *Sliding Mode Control: Theory Appl.* Crc Press, 1998.