

# Physics-Informed Neural Network for Multirotor Slung Load Systems Modeling

Gil Serrano\*, Marcelo Jacinto\*, José Ribeiro-Gomes\*, João Pinto\*,  
Bruno J. Guerreiro, Alexandre Bernardino, and Rita Cunha

**Abstract**—Recent advances in aerial robotics have enabled the use of multirotor vehicles for autonomous payload transportation. Resorting only to classical methods to reliably model a quadrotor carrying a cable-slung load poses significant challenges. On the other hand, purely data-driven learning methods do not comply by design with the problem’s physical constraints, especially in states that are not densely represented in training data. In this work, we explore the use of physics-informed neural networks to learn an end-to-end model of the multirotor-slung-load system and, at a given time, estimate a sequence of the future system states. An LSTM encoder-decoder with an attention mechanism is used to capture the dynamics of the system. To guarantee the cohesiveness between the multiple predicted states of the system, we propose the use of a physics-based term in the loss function, which includes a discretized physical model derived from first principles together with slack variables that allow for a small mismatch between expected and predicted values. To train the model, a dataset using a real-world quadrotor carrying a slung load was curated and is made available. Prediction results are presented and corroborate the feasibility of the approach. The proposed method outperforms both the first principles physical model and a comparable neural network model trained without the physics regularization proposed.

**Code + Dataset:** <https://github.com/GilSerrano/pinn-air>

**Video:** [https://youtu.be/4N\\_nnaGsUOg](https://youtu.be/4N_nnaGsUOg)

## I. INTRODUCTION

System identification is a fundamental task in model-based control of dynamical systems, as it involves the estimation of the underlying dynamics from input-output data – due to either the model being hard to derive or its parameters being unknown *a priori*. Characterizing a physical system through observed data is critical to understanding its dynamics and allowing the design of controllers that enable the system to follow reference signals with a certain desired behavior.

\*These authors contributed equally.

The work of G. Serrano, M. Jacinto, J. Ribeiro-Gomes, and J. Pinto was supported by the PhD Grants PRT/BD/154275/2022, 2022.09587.BD, 2022.12145.BD from MIT Portugal and Fundação para a Ciência e a Tecnologia (FCT), Portugal. This work was also supported by FCT, Portugal through LARSyS [DOI: 10.54499/LA/P/0083/2020, 10.54499/UIDP/50009/2020, and 10.54499/UIDB/50009/2020], project CAPTURE [DOI: 10.54499/PTDC/EEI-AUT/1732/2020], and project HAVATAR [DOI: 10.54499/PTDC/EEI-ROB/1155/2020].

The authors are with the Institute for Systems and Robotics (ISR), Laboratory of Robotics and Engineering Systems (LARSyS), Instituto Superior Técnico, University of Lisbon, Portugal. B. J. Guerreiro is also with CTS/Uninova, NOVA School of Science and Technology (FCT-NOVA), Caparica, Portugal. E-mails: {gil.serrano, josepgomes}@tecnico.ulisboa.pt, {mjacinto, jpinto, alex, rita}@isr.tecnico.ulisboa.pt, bj.guerreiro@fct.unl.pt

Classical approaches to system identification mainly rely on the statistical analysis of the input-output relation [1]–[3] to extract crucial information from collected data. Time and frequency-domain analyses are some of the more common approaches. Using a pseudo-random binary signal (PBRS) that excites most frequencies, one can, with relative ease, design linear auto-regressive models that capture the most predominant modes of the system. In contrast, neural networks with nonlinear activations are able to learn nonlinear approximations of complex system dynamics.

Recent and promising approaches have attempted to use neural networks to model unknown system dynamics. Physics-Informed Neural Networks (PINNs [4]) are neural networks that are trained to solve supervised learning tasks while taking into account the physical constraints of the world. By exploiting the knowledge of the physical laws that govern the system under analysis in the learning process, PINNs overcome the low data availability that may occur when dealing with real systems. This prior knowledge acts as a regularization term that limits the space of admissible solutions. Other approaches resort to a combination of Recursive Neural Networks (RNNs) with variational auto-encoders [5] or hybrid physical models [6], [7] to capture the recursive nature of dynamical systems.

The combination of data-driven methods with physically plausible models is specially important in systems with large degrees of uncertainty. For instance, the interaction between a multirotor aircraft and a slung load (Fig. 1) makes system identification challenging, due to phenomena such as downwash on the load. While this system is relatively simple, it has interesting use cases such as aerial transportation [8]. Classical approaches for system identification may not be well-suited for this problem, as they struggle to accurately capture the nonlinear dynamics of the system over long time periods, since the maneuvers that can be performed to excite the modes of the coupled system (drone and slung load) are limited. Deep learning approaches, on the other hand, have shown promise in deducing complex dynamics and how they evolve over time [4], [5], [9]–[11]. The development of controllers that leverage learned dynamics could be a significant step toward the successful operation of multirotor systems carrying slung loads.

**Contributions:** This work proposes the use of a Long-Short Term Memory (LSTM) encoder-decoder network with an attention mechanism to model end-to-end the system dynamics and perform multi-step predictions of a quadrotor carrying a slung load. The main contributions are:



Fig. 1: Intel Aero quadrotor carrying a cube-shaped slung load attached by a string in a dedicated motion capture arena.

- an encoder-decoder architecture to solve the prediction problem making use of slack variables and a physics-based loss term – acting as a regularization term that limits the space of search and designed by taking into account a discrete approximation of the system derived from first principles;
- a novel real-world dataset collected in a motion capture (MOCAP) arena that was obtained by conducting 48 independent flight maneuvers.

The paper is organized as follows: Section II summarizes the research conducted on the intersection of deep learning and system identification for the control of dynamical systems; Section III presents the problem formulation, with a focus on the multirotor dynamics, and the proposed network architecture and loss function; Section IV details the data acquisition and pre-processing required to then train the network, the evaluation metrics employed, and presents the results obtained. Concluding remarks are offered in Section V along with a future line of research.

## II. RELATED WORK

Recent approaches have attempted to use neural networks to model unknown system dynamics. One such approach is the Sparse Identification of Nonlinear Dynamics (SINDy) [12] framework, which proposes expressing the dynamics, given by ordinary differential equations (ODEs), as a linear combination of nonlinear basis functions. To determine the contribution of each basis function, a convex sparse linear regression problem is formulated. An extension of SINDy proposes combining this idea with the use of a neural network with an autoencoder structure to simultaneously discover a low-dimensional latent space of coordinates and functions that model the dynamics of the system [9].

Another concept in system identification is Physics-Informed Neural Networks (PINNs) [4]. These networks are capable of encoding the fundamental dynamics that govern a dataset, given measurements of the evolution of the system described by a nonlinear partial differential equation (PDE)  $\partial_t g + \mathcal{N}[g; \lambda] = 0$ , where  $g(x, t)$  denotes the latent solution, dependent on variables  $x$  and  $t$ , and  $\mathcal{N}[g; \lambda]$  is a nonlinear differential operator parameterized by a vector  $\lambda$ . By minimizing the sum of the MSE (Mean Squared

Error) of the boundary conditions and the MSE of the fit of a set of points in  $(x, t)$  with regards to the differential equation, PINNs are able to estimate the parameters  $\lambda$  that best describe the observed data. Standard PINN models are useful in a variety of problems, however, they hinge on the knowledge of  $\mathcal{N}$  and are not suited for control tasks, as they are not designed to handle variable control inputs.

Some work has been done to make PINNs useful in control applications. In [10], Physics-Informed Neural Nets for Control (PINCs) were introduced, which extend PINNs by adding control actions as inputs. PINCs are able to identify a dynamical system from data and its governing ODEs, and can also simulate differentiable equations faster than conventional numerical methods. This approach is able to propagate the equations for variable horizons without serious degradation in the predictions, measured with absolute error, which is critical for model predictive control applications.

The research presented in [11] also extended PINNs by adding control actions as network inputs. With these modifications, the authors were able to combine PINNs with Nonlinear Model Predictive Control (NMPC). At each instant, the neural network receives the current time step, state, and input and predicts the next state, thus effectively replacing an explicit model in the NMPC algorithm. Simulations demonstrate that this approach can solve a tracking problem for a robotic manipulator, surpassing classical time-integration methods in terms of computational efficiency. The performance is evaluated with the Root Mean Squared Error (RMSE) along the reference trajectory.

In [5], Deep State Space (DSS) models are presented as a system identification approach that can describe a wide range of dynamics due to the flexibility of deep neural networks. The focus of that work is on sequential data and the proposed strategies are based on RNNs, as they are useful in modeling sequences of variable length, and variational autoencoders.

The approach presented in this work is inspired by the original work on PINNs [4] and DSS [5] with the key difference being the combination of RNNs with an encoder-decoder structure to model a discretized dynamical system with variable control inputs, and the use of a physical model derived from first principles applied to the network output

during training to restrict the space of admissible predictions, combined with auxiliary slack variables that enhance the ability to account for unmodeled system dynamics.

### III. PROPOSED METHOD

The dynamics of a quadrotor carrying a slung load are governed by a differential equation that can be derived as in Yu et al. [13], where it is assumed that the load is rigidly connected to the center of mass of the quadrotor by a fixed-length cable that is taut during flight. The dynamical model adopted does not account for aerodynamic drag, the downwash on the load, structure bending modes, nor external disturbances, which can prove hard to quantify accurately in practice. In this section, we present the physical model in its discrete form. The network that is described in Section III-B uses this knowledge during training to learn the dynamics of the full system composed of the quadrotor and a slung load.

#### A. System model

Let  $\mathcal{W}$  denote the world reference frame, which corresponds to the canonical basis of  $\mathbb{R}^3$ ,  $\mathcal{W} = \{\mathbf{e}_i\}_{i=1,2,3}$ , and  $\mathcal{B}$ , the body frame attached to the vehicle with origin at the center of mass of the quadrotor. The position  $\mathbf{p}^L$  and velocity  $\mathbf{v}^L$  of the payload in  $\mathcal{W}$ , considering a sampling time  $h \in \mathbb{R}^+$ , at time instant  $k+1$  are approximated by

$$\begin{cases} \mathbf{p}_{k+1}^L &= \mathbf{p}_k^L + h\mathbf{v}_k^L + \frac{1}{2}h^2\mathbf{u}_k^* \\ \mathbf{v}_{k+1}^L &= \mathbf{v}_k^L + h\mathbf{u}_k^* \end{cases}, \quad (1)$$

with the virtual input  $\mathbf{u}_k^*$  given by

$$\mathbf{u}_k^* = -\frac{1}{m_L}T_k^L\boldsymbol{\xi}_k + g\mathbf{e}_3, \quad (2)$$

where  $m_L \in \mathbb{R}^+$  denotes the mass of the load,  $T_k^L \in \mathbb{R}$  is the tension in the cable,  $\boldsymbol{\xi}_k \in \mathbb{S}^2$  is the load cable direction, and  $g = 9.8 \text{ m s}^{-2}$  denotes Earth's gravity. The evolution of the position  $\mathbf{p}$  and velocity  $\mathbf{v}$  in  $\mathcal{W}$  of the multirotor are provided by

$$\begin{cases} \mathbf{p}_{k+1} &= \mathbf{p}_k + h\mathbf{v}_k + \frac{1}{2}h^2(\bar{\mathbf{u}}_k + \mathbf{u}_k^*) \\ \mathbf{v}_{k+1} &= \mathbf{v}_k + h(\bar{\mathbf{u}}_k + \mathbf{u}_k^*) \end{cases}, \quad (3)$$

where the virtual input  $\bar{\mathbf{u}}_k$  is given by

$$\bar{\mathbf{u}}_k = \frac{T_k^Q}{m_Q}\mathbf{R}_k\mathbf{e}_3 - \frac{m_Q + m_L}{m_Q}g\mathbf{e}_3, \quad (4)$$

where  $m_Q \in \mathbb{R}^+$  denotes the mass of the vehicle,  $T_k^Q \in \mathbb{R}^+$  corresponds to the total thrust produced by the rotors, and  $\mathbf{R}_k \in \text{SO}(3)$  corresponds to the vehicle orientation. The orientation kinematics are propagated as follows,

$$\mathbf{R}_{k+1} = \mathbf{R}_k \exp\{h\mathbf{S}(\boldsymbol{\omega}_k)\}, \quad (5)$$

with  $\boldsymbol{\omega}_k \in \mathbb{R}^3$  representing the quadrotor angular velocity expressed in the body frame  $\mathcal{B}$ , and  $\mathbf{S}(\cdot)$  is an isomorphism between  $\mathbb{R}^3$  and the space of  $3 \times 3$  skew-symmetric matrices defined by the condition  $\mathbf{S}(\mathbf{a})\mathbf{b} = \mathbf{a} \times \mathbf{b}$ ,  $\forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^3$ . The load cable direction vector can be obtained by

$$\boldsymbol{\xi}_k = \frac{\mathbf{p}_k^L - \mathbf{p}_k}{\ell}, \quad (6)$$

where  $\ell$  denotes the cable length. The tension on the cable can be computed according to

$$T_k^L = -\frac{m_L}{m_L + m_Q}\boldsymbol{\xi}_k^\top \mathbf{R}_k T_k^Q \mathbf{e}_3 + \frac{m_L m_Q}{m_L + m_Q} \ell \left\| \boldsymbol{\Omega}_k^L \right\|^2, \quad (7)$$

where  $\boldsymbol{\Omega}_k^L \in \mathbb{R}^3$  is the angular velocity of the load and its temporal evolution is determined by

$$\boldsymbol{\Omega}_{k+1}^L = (1-d)\boldsymbol{\Omega}_k^L - \frac{h}{m_Q \ell} \mathbf{S}(\boldsymbol{\xi}_k) \mathbf{R}_k T_k^Q \mathbf{e}_3, \quad (8)$$

with  $d \in [0, 1]$ , a damping term. The observed variables of the system state are stored on the tuple  $\mathbf{x}_k = \{\mathbf{p}_k, \mathbf{v}_k, \mathbf{q}_k, \mathbf{p}_k^L\}$ , corresponding to the position, linear velocity, and orientation – expressed using unit quaternions  $\mathbf{q}_k \in \mathbb{S}^3$  – of the vehicle, and the position of the load  $\mathbf{p}_k^L$  expressed in  $\mathcal{W}$ . The term  $\boldsymbol{\Omega}_k^L$  is not included in  $\mathbf{x}_k$ , as it is not measured directly. The orientation is represented as a unit quaternion because this parameterization has been shown to converge faster and produce better results than rotation matrices in the context of neural networks [14]. The control input of the complete system is given by  $\mathbf{u}_k = \{T_k^Q, \boldsymbol{\omega}_k\}$  and is comprised of the total thrust produced by the quadrotor and its body angular rates. The evolution of  $\mathbf{x}_k$  is, therefore, described by  $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k, \boldsymbol{\Omega}_k^L)$ .

#### B. Network architecture

Given a time series of observations of state variables and control inputs  $\{\mathbf{x}_k, \mathbf{u}_k\}$  up to time step  $M-1$ , the network aims to recursively predict the evolution of the state of the quadrotor-slung-load system for the next  $N$  time steps, given only the control input that is applied to the vehicle at those instants. The architecture of the proposed network is shown in Fig. 2. We start by providing  $\{\mathbf{x}_k, \mathbf{u}_k\}$  to an encoder – these vectors are passed through fully connected layers with GELU activations to discover a latent space representation of the input data  $\{\mathbf{z}_k\}$ ,  $k = 0, 1, \dots, M-1$ . This sequence of new feature vectors is fed to a multilayer LSTM to produce a hidden state  $\mathbf{h}_M$  and the encoder output  $\mathbf{v}_{\text{enc}}$ .

To obtain the estimates of the system state  $\hat{\mathbf{x}}_k$  in the subsequent time steps, we use a decoder composed of a multilayer LSTM network with a Bahdanau attention mechanism [15]. At  $k = M$ , the decoder receives the encoder output  $\mathbf{v}_{\text{enc}}$ , the previous hidden state  $\mathbf{h}_M$ , and the last known system state and the actual control input  $\{\mathbf{x}_{M-1}, \mathbf{u}_{M-1}\}$ . The system state and the control input are fed to fully connected layers and the output  $\mathbf{z}_{M-1}$  is passed to the attention layer, along with the encoder output and previous hidden state. Then, the context  $\mathbf{c}$  that is calculated by the attention mechanism is concatenated with  $\mathbf{z}_{M-1}$  and fed to the LSTM. The output of the LSTM is then passed through fully connected layers that convert it to the original state space, thus obtaining the estimate  $\hat{\mathbf{x}}_M$ , which is concatenated with the next actual input to make a recursive prediction. Additionally, a set of slack variables  $\mathbf{s}_M$  that is used by the loss function to impose soft-physical constraints on the system predictions during training is generated. In the following time steps,  $k = M+1, \dots, M+N-1$ , the previous system state estimate

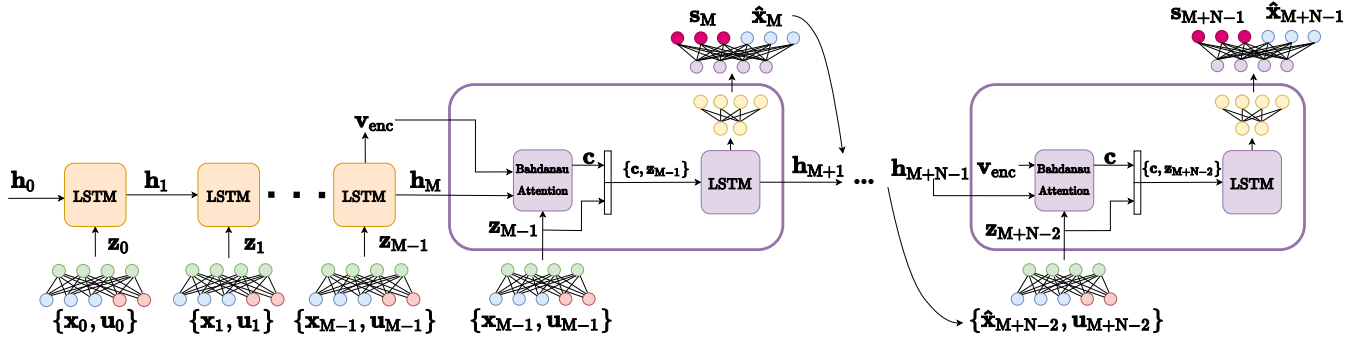


Fig. 2: Proposed encoder-decoder architecture with an attention mechanism. The network encoder receives a temporal history of  $M$  system states and control inputs. The decoder receives the hidden state and output of the encoder, the last measured state of the system  $\mathbf{x}_{M-1}$ , and the previous input of the system  $\mathbf{u}_{M-1}$ , which are combined in an attention module to produce a prediction of the next system state  $\hat{\mathbf{x}}_M$  and the corresponding slack variables  $\mathbf{s}_M$ . Each prediction of the decoder module is concatenated with an input control signal to generate multi-step predictions of the system state in the following iterations.

$\hat{\mathbf{x}}_{k-1}$  is fed to the fully connected layers, instead of the actual system state, along with the control input.

To train the model, a loss function given by the sum of four separate terms is considered

$$\mathcal{L} = \mathcal{L}_{\text{fit}} + \mathcal{L}_{\text{physics}} + \mathcal{L}_{\text{projection}} + \mathcal{L}_{\text{slack}}. \quad (9)$$

The first term of the loss is included to enforce the network to fit the training data,

$$\begin{aligned} \mathcal{L}_{\text{fit}} = & \sum_{n=M}^{M+N-1} \lambda_1 \|\hat{\mathbf{p}}_n - \mathbf{p}_n\|_{\mathbf{M}_n}^2 + \lambda_2 \|\hat{\mathbf{v}}_n - \mathbf{v}_n\|_{\mathbf{M}_n}^2 \\ & + \lambda_3 \|\mathbf{e}(\hat{\mathbf{q}}_n, \mathbf{q}_n)\|_{\mathbf{M}_n}^2 + \lambda_4 \|\hat{\mathbf{p}}_n^L - \mathbf{p}_n^L\|_{\mathbf{M}_n}^2, \end{aligned} \quad (10)$$

where  $\forall \mathbf{y} \in \mathbb{R}^m$ ,  $\|\mathbf{y}\|_{\mathbf{M}_n}^2 := \mathbf{y}^\top \mathbf{M}_n \mathbf{y}$ ,  $\mathbf{M}_n := \mathbf{I} e^{-\alpha(n-M)}$ , with  $\alpha > 0$  and  $\mathbf{I}$  being the identity matrix of appropriate dimension; the map  $\mathbf{e} : \mathbb{S}^3 \times \mathbb{S}^3 \rightarrow \mathbb{R}^4$  is defined as

$$\mathbf{e}(\mathbf{q}_1, \mathbf{q}_2) = \mathbf{q}_1 \otimes \mathbf{q}_2^{-1} - \mathbf{q}_{\text{id}}, \quad (11)$$

where  $\mathbf{q}_{\text{id}}$  denotes the identity quaternion, to measure the error between two quaternions representing the orientation of the vehicle. The second term,

$$\begin{aligned} \mathcal{L}_{\text{physics}} = & \phi \left\| \hat{\mathbf{x}}_M - f(\mathbf{x}_{M-1}, \mathbf{u}_{M-1}, \Omega_{M-1}^L) + \mathbf{s}_M \right\|_{\mathbf{Q}_M}^2 \\ & + \sum_{n=M}^{M+N-2} \phi \left\| \hat{\mathbf{x}}_{n+1} - f(\hat{\mathbf{x}}_n, \mathbf{u}_n, \Omega_n^L) + \mathbf{s}_{n+1} \right\|_{\mathbf{Q}_{n+1}}^2, \end{aligned} \quad (12)$$

ensures the system dynamics restrict the admissible search space. The function  $f$  represents the discretized drone dynamics introduced in Section III-A so that  $\hat{\mathbf{x}}_{k+1} \approx f(\hat{\mathbf{x}}_k, \mathbf{u}_k, \Omega_k^L)$ . The term  $\mathcal{L}_{\text{physics}}$  acts, therefore, as a regularization mechanism that penalizes solutions that do not satisfy the system dynamics. However, since the physical model is not ideal, there are discrepancies between the discretized model predictions and observations gathered from the real system, causing the estimates to diverge increasingly as prediction horizons become longer. As such, the slack variables  $\mathbf{s}_k$  are introduced as extra states predicted by the network<sup>1</sup>, which allow for a small mismatch between the

<sup>1</sup>The slack variables are only used during the training phase, and their value is discarded at evaluation time.

predictions of the physical model and the network. Moreover, in the penalty term, a sequence of diagonal matrices  $\mathbf{Q}_n := \mathbf{I} e^{-\beta(n-M)}$ ,  $\beta > 0$ , that converge exponentially to the zero matrix is introduced to penalize the mismatch between estimates and model predictions in the first instants and gradually less further into the future.

To represent a rotation, a quaternion must have unit norm and the network should estimate it as such. The loss function incorporates, therefore, this constraint as follows,

$$\mathcal{L}_{\text{projection}} = \psi \sum_{n=M}^{M+N-1} |1 - \|\hat{\mathbf{q}}_n\||^2. \quad (13)$$

Finally, it is desirable that the slack variables  $\mathbf{s}_k$  are close to zero so that the network predictions follow the physical model as closely as possible. To explicitly account for such requirement, the last term of the loss function is given by

$$\mathcal{L}_{\text{slack}} = \rho \sum_{n=M}^{M+N-1} \|\mathbf{s}_n\|_{\mathbf{Q}_n}^2. \quad (14)$$

The hyperparameters  $\lambda_1$ ,  $\lambda_2$ ,  $\lambda_3$ ,  $\lambda_4$  in (10);  $\phi$  in (12);  $\psi$  in (13); and  $\rho$  in (14) determine the relative weighting of each of the squared error terms in  $\mathcal{L}_{\text{fit}}$ ,  $\mathcal{L}_{\text{physics}}$ ,  $\mathcal{L}_{\text{projection}}$ , and  $\mathcal{L}_{\text{slack}}$ , respectively.

## IV. RESULTS

In this section, we describe the data acquisition and processing procedures. Experiments are conducted to assess the performance of the proposed model and training strategy.

### A. Data Acquisition and Model Training

To acquire real-world data, a 100 g cube with 10 cm edges was attached to an Intel Aero RTF drone with a 60 cm string (Fig. 1). The drone was commanded to perform various flight maneuvers to excite different modes of the system, as depicted in Fig. 3. The control signal  $\mathbf{u}$  was determined by an offboard controller. The acquired data containing the state and input of the system was synchronized and sampled at  $h = 30$  ms. A total of 48 different experiments was conducted, with duration varying from 197 to 799 time steps.

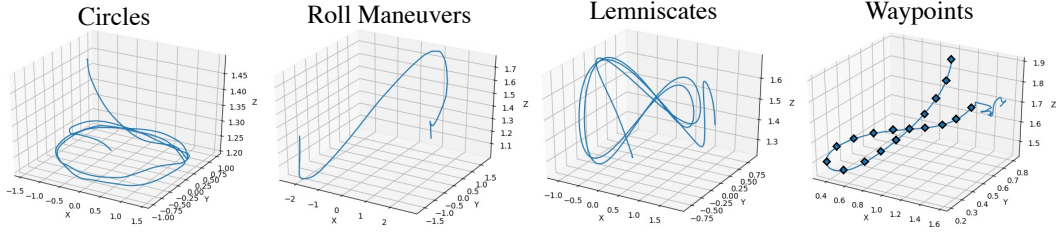


Fig. 3: Subset of the trajectories executed by the quadrotor carrying a slung load during data acquisition.

Each time series was divided into chunks of a fixed duration  $M = 50$  time steps, containing the state and control inputs to the system, and the model is trained to predict the state of the system for the next  $N = 25$  time steps. Additional evaluation tests were also conducted considering  $N = 50$ .

A sliding window approach with stride 1 was used to generate sequences from the flight data recorded. A split of 58-17-25 was adopted, corresponding to 9,959 sequences for training, 2,837 for validation, and 4,252 for testing. Each model is trained using Adam [16] with a learning rate of  $1 \times 10^{-4}$  and an  $\ell_2$  weight decay of 0.05. The training procedure is performed for 3,000 epochs, with early-stopping and batches composed of 256 sequences on an NVIDIA RTX 3090 GPU, taking 6 s per epoch, and a total of 5 h per model trained. The hyperparameters  $\lambda_1 = \lambda_2 = \lambda_3 = \lambda_4 = 20$ ,  $\phi = 5$ ,  $\psi = 10$ ,  $\rho = 0.1$ ,  $\alpha = 0.1$ ,  $\beta = 0.6$  were considered, and the trained model has 238,906 parameters.

### B. Performance Metrics

To assess the performance of the proposed model and training strategy, we compare our model with a physics-based open-loop state predictor, derived from first principles, i.e. the model in Section III-A. Additionally, a performance comparison with a baseline LSTM encoder-decoder network trained on the same training data, but without slack variables and without the  $\mathcal{L}_{\text{physics}}$  and  $\mathcal{L}_{\text{slack}}$  loss terms is conducted – enabling the evaluation of the impact of the physics constraints on the performance of the model. The performance of each model is analyzed by computing the evolution of the Mean Average Error (MAE) over all the test set sequences at each timestep  $k$  according to

$$\text{MAE}_k = \frac{1}{\beta} \sum_{j=1}^{\beta} \left\| \zeta_k^j - \hat{\zeta}_k^j \right\|, \quad (15)$$

where  $\zeta_k^j$  and  $\hat{\zeta}_k^j$  are a subset of the groundtruth and prediction of the system state for the test sample  $j$ , respectively, and  $\beta$  is the number of test sequences. Note that the quaternion error is computed according to the formula presented in (11). The RMSE is additionally computed for each model over the entire test set.

### C. Prediction Performance

In Fig. 4 the evolution of the MAE over time for the proposed neural network with the discretized physical model as well as the Q1 and Q3 quartiles of the quadrotor position, velocity, orientation, and payload position are shown. From

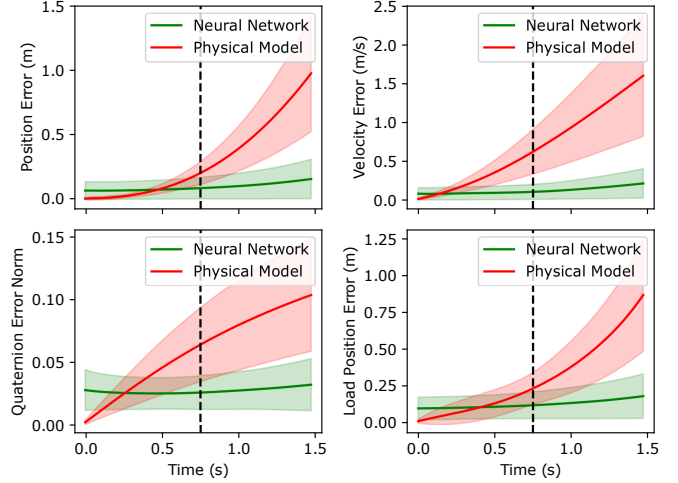


Fig. 4: Comparison between the evolution of the MAE and standard deviation of the prediction error of the proposed model and the physics-based state predictor, evaluated on the test set. The vertical dashed line denotes the prediction horizon of  $N = 25$  time steps (750 ms) used during training.

these results, it is possible to conclude that in the initial prediction steps, the error is smaller for the physics-based predictor, but as the prediction time increases, the neural network outperforms the classical model and its standard deviation increases at a slower rate. The dashed vertical line denotes the 25<sup>th</sup> step of the prediction horizon adopted during training. The plots also demonstrate that even though the model is only trained to make predictions for 750 ms, it can be used over longer prediction horizons at evaluation time. In Fig. 5 we showcase the results of the recursive prediction of the state of the system for a single sequence of the test set. The first  $M = 50$  time steps (1.5 s) are used as the initial input to the neural network, and the output of the proposed model is depicted in dashed lines for the following timesteps over the groundtruth. The time series predicted by the discretized dynamical model are depicted in dotted lines. In this particular example, it is possible to conclude from visual inspection that the learned model can accurately follow the evolution of each state during the first  $N = 25$  prediction steps, while the predictions generated by the dynamical model start to deviate from the groundtruth, particularly the velocity predictions. This is due to the fact that the discretized dynamical model does not account for disturbances such as aerodynamic drag and the downwash on the slung load, in contrast to the network which is able

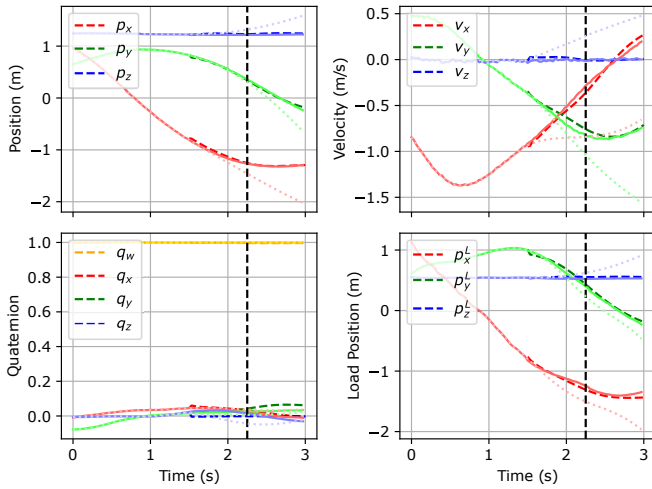


Fig. 5: Prediction of the vehicle position, linear velocity, orientation, and the position of the slung load as a function of time for a single sequence of the test set. The time series predicted by the network are depicted in **dashed lines** over the expected values in **solid lines**. The time series predicted by the discretized dynamics model are depicted in **dotted lines** for comparison. The vertical dashed lines denote the prediction horizon of 750 ms used during training.

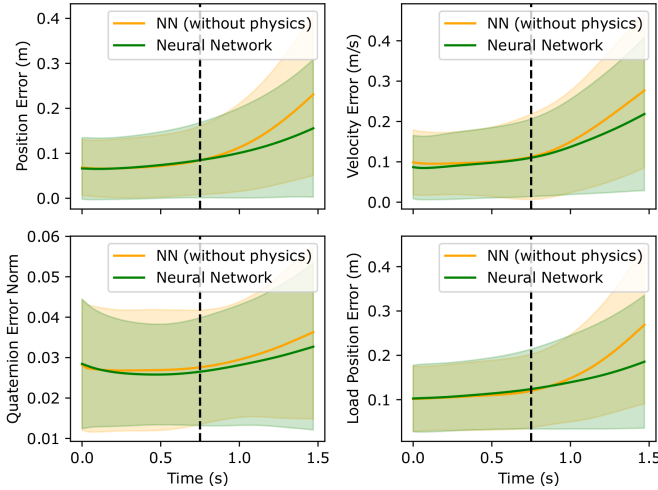


Fig. 6: Evolution of the MAE and standard deviation of the prediction error of the proposed model with a baseline model trained without physics constraints. The vertical dashed line denotes the prediction horizon used during training.

to learn the unmodeled dynamics.

To evaluate the impact of imposing physical constraints on the loss function during training, Fig. 6 compares the evolution of the MAE of the proposed model with an equivalent baseline model trained without the physical constraints. From these results, we confirm that enforcing physical constraints on the model reduces the mean prediction error over a longer prediction horizon, although the gains over the baseline neural network model are also shown to be smaller in comparison to the gains obtained over the physical model predictor. The RMSE results presented in Table I obtained over the entire test set, for the discretized physics model, and

TABLE I: RMSE of the different models, over the entire test set for a prediction horizon of  $N = 50$  time steps.

Model	RMSE				
	Position (m)	Velocity (m/s)	Quaternion	Payload Position (m)	Combined
Physics	0.458	0.924	0.073	0.424	1.879
Proposed Model (w/o physics)	0.153	0.193	0.033	0.182	0.561
Proposed Model (w/o slack vars)	0.153	0.200	0.034	0.184	0.571
Proposed Model	<b>0.136</b>	<b>0.175</b>	<b>0.032</b>	<b>0.166</b>	<b>0.508</b>

TABLE II: Inference time.

Horizon		RTX 3090		Jetson Orin Nano	
$M$	$N$	mean (ms)	std dev (ms)	mean (ms)	std dev (ms)
50	25	10.76	0.53	59.44	0.26
50	50	21.24	0.99	115.26	0.59

for the neural network trained with and without the physics loss term and slack variables further support the analysis. The inference times for different horizons  $M$  and  $N$  on an NVIDIA RTX 3090 GPU and on an NVIDIA Jetson Orin Nano are presented in Table II and support the feasibility of the proposed model.

## V. CONCLUSION

We have proposed a network architecture to recursively predict the evolution of a quadrotor-slung-load system from observations of its state and control inputs. The network features an encoder-decoder structure with attention that provides estimates of the state of the quadrotor-slung-load system in subsequent time steps along with auxiliary slack variables that allow to better account for unmodeled system dynamics. Moreover, the modeled quadrotor-slung-load dynamics discounted over the prediction horizon are incorporated into the loss function to train the network so that more accurate predictions of the evolution of the coupled system are obtained.

The network was trained using data from various flight maneuvers performed in a controlled experimental setup. The results show the model is able to accurately predict how the dynamics of the system evolve in time, outperforming the predictions made by a model derived from first principles. The impact of the physics constraints and slack variables in the loss function is assessed and it is verified that combining these terms together results in a significant reduction of the prediction error. Moreover, the proposed model is able to capture the non-trivial dynamics of the quadrotor-slung-load system and generalizes well for a longer prediction horizon than the one considered during training. The inference times also corroborate the feasibility of the approach in a practical setting for different horizons.

Further research includes expanding this approach to a broader range of vehicles and dynamic systems. The design of model-based control strategies that leverage the proposed prediction technique is a relevant line of research to pursue, with special emphasis on long prediction horizon scenarios.

## ACKNOWLEDGMENT

The authors gratefully acknowledge Chrysoula Zerva and André F. T. Martins for their suggestions to improve the quality of this work.

## REFERENCES

- [1] L. Ljung, *System Identification: Theory for the User*, ser. Prentice Hall Information and System Sciences Series. Prentice Hall PTR, 1999.
- [2] T. Söderström and P. Stoica, *System identification*. Prentice-Hall International, 1989.
- [3] P. Van Overschee and B. De Moor, *Subspace identification for linear systems: Theory—Implementation—Applications*. Springer Science & Business Media, 2012.
- [4] M. Raissi, P. Perdikaris, and G. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [5] D. Gedon, N. Wahlström, T. B. Schön, and L. Ljung, “Deep state space models for nonlinear system identification,” *IFAC-PapersOnLine*, vol. 54, no. 7, pp. 481–486, 2021, 19th IFAC Symposium on System Identification SYSID 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405896321011800>
- [6] N. Mohajerin, M. Mozifian, and S. Waslander, “Deep learning a quadrotor dynamic model for multi-step prediction,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2454–2459.
- [7] P. Karle, F. Török, M. Geisslinger, and M. Lienkamp, “Mixnet: Physics constrained deep neural motion prediction for autonomous racing,” *IEEE Access*, pp. 1–1, 2023.
- [8] X. Liang, H. Yu, Z. Zhang, Y. Wang, N. Sun, and Y. Fang, “Unmanned quadrotor transportation systems with payload hoisting/lowering: Dynamics modeling and controller design,” in *2020 5th International Conference on Advanced Robotics and Mechatronics (ICARM)*, 2020, pp. 666–671.
- [9] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, “Data-driven discovery of coordinates and governing equations,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 45, pp. 22 445–22 451, 2019.
- [10] E. A. Antonelo, E. Camponogara, L. O. Seman, E. R. de Souza, J. P. Jordanou, and J. F. Hübner, “Physics-informed neural nets-based control,” *CoRR*, vol. abs/2104.02556, 2021. [Online]. Available: <https://arxiv.org/abs/2104.02556>
- [11] J. Nicodemus, J. Kneifl, J. Fehr, and B. Unger, “Physics-informed neural networks-based model predictive control for multi-link manipulators,” *IFAC-PapersOnLine*, vol. 55, no. 20, pp. 331–336, 2022, 10th Vienna International Conference on Mathematical Modelling MATHMOD 2022.
- [12] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [13] G. Yu, D. Cabecinhas, R. Cunha, and C. Silvestre, “Nonlinear backstepping control of a quadrotor-slung load system,” *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 5, pp. 2304–2315, 2019.
- [14] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, “On the continuity of rotation representations in neural networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5745–5753.
- [15] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: <http://arxiv.org/abs/1409.0473>
- [16] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>