2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Credible Online Dynamics Learning for Hybrid UAVs

David Rohr[1], Nicholas Lawrance[1,2], Olov Andersson[1] and Roland Siegwart[1]

***Abstract*— Hybrid unmanned aerial vehicles (H-UAVs) are highly versatile platforms with the ability to transition between rotary- and fixed-wing flight. However, their (aero)dynamics tend to be highly nonlinear which increases the risk of introducing safety-critical modeling errors in a controller. Designing a safe, yet not too cautious controller, requires a *credible* model which provides accurate dynamics uncertainty quantification. We present a data-efficient, probabilistic semi-parametric dynamics modeling approach that allows for online, filter-based inference. The proposed model leverages prior knowledge using a nominal parametric model, and combines it with residuals in the form of sparse Gaussian processes to account for possibly unmodeled forces and moments. Uncertain nominal and residual parameters are jointly estimated using Bayesian filtering. The resulting model accuracy and the reliability of its predicted uncertainty are analyzed for both a simulated and a real example, where we learn the 6DoF nonlinear dynamics of a tiltwing H-UAV from a few minutes of flight data. Compared to a residual-free nominal model, the proposed semi-parametric approach provides increased model accuracy in relevant parts of the flight envelope and substantially higher credibility overall.**

I. Introduction

H-UAVs are versatile aircraft that transition between fixed-wing and rotary-wing flight modes and thus feature both efficient and fast cruise flight as well as vertical take-off/landing (VTOL) capabilities. The large flight envelope of H-UAVs enables many attractive applications such as long-range point-to-point parcel delivery. However, it often entails highly nonlinear and vastly different flight dynamics from hover to cruise. This considerably complicates accurate modeling and increases the risk of creating a model with errors that degrade the performance of model-based controllers or, more critically, impair controller stability and thereby system safety. The latter risk is further stressed considering that the data used to train and/or improve a model might only be available through actual flight tests, e.g., when wind-tunnel testing is infeasible due to costs or does not provide the required range of data. As such, the UAV needs to be controlled using a model which is initially uncertain. To still enable safe data collection for model improvement—and safe flight in general—it is therefore key that a dynamics model also provides reliable uncertainty quantification which can be used in a probabilistically safe control framework. Furthermore, an accurate indication of model uncertainty can be leveraged for targeted, active exploration of the dynamics [1]–[3], which is highly desirable considering the limited flight time available in practice.

[1] Autonomous Systems Lab, ETH Zurich, Leonhardstrasse 21, 8092 Zurich, Switzerland, <firstname.lastname>@mavt.ethz.ch
[2] Robotics and Autonomous Systems Group, CSIRO Data61, QLD 4069, Australia, nicholas.lawrance@csiro.au

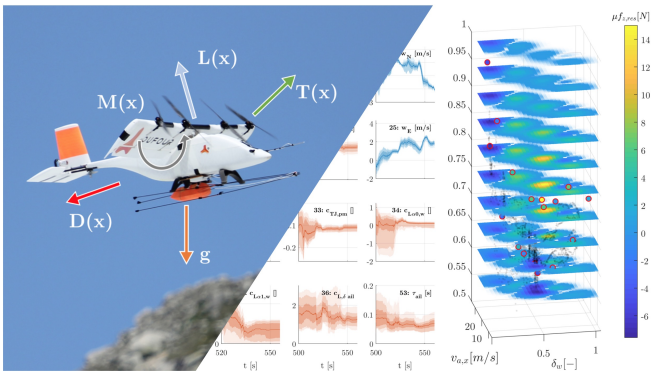

Fig. 1. The presented online semi-parametric learning framework estimates the nonlinear (aero)dynamics of H-UAVs from flight data and provides reliable model uncertainty quantification. Dynamic states, model parameters and residual terms based on sparse Gaussian processes (SPGP) are jointly estimated in real-time using filter-based inference. For qualitative illustration, we show a collection of mean estimates for states (blue), system parameters (red), and a residual component $\mathbb{R}^3 \rightarrow \mathbb{R}$ in the form of slices (right).

Ultimately, we aim to develop *safe, active learning* approaches to create accurate models online during flight. The work presented in this paper is a critical component to reach this goal: a dynamics modeling and learning framework that provides:

- *Expressiveness* for high accuracy and model-based control.
- *Credibility* for accurate uncertainty quantification to encourage probabilistically safe, yet not too cautious, control and exploration.
- *Sample efficiency* to learn from limited real-world data.
- *Real-time learning* for online model adaptation and prompt (re-)planning of exploration trajectories in flight.

1) Related Work: Parametric *grey-box* approaches based on domain knowledge and first principles form a strong prior for learning dynamics models, including H-UAVs [4]–[6]. The informed structure and embedded prior knowledge allow for fast and sample-efficient learning, and offer good generalization. However, we cannot generally assume that these models can capture all nuances relevant to a particular application. For example, the model structure presented in [6] fails to explain the pitch dynamics of the H-UAV in Fig. 1 during backward flight, predicting available control margins when, in fact, the controller saturates the system.

To improve accuracy, so-called *semi-parametric* modeling approaches combine prior-knowledge based parametric (hereafter *nominal*) components for sample efficiency with general-purpose, data-driven *residual* terms to compensate for remaining errors. Successful applications are demonstrated on several different robotics platforms, including

unmanned aerial vehicles (UAVs), with popular residual terms including neural networks [7], [8], mixture models [9] and full- or sparse Gaussian processs (GPs) [10]–[14]. However, while these works achieve potentially high accuracy, they only provide a single (e.g. maximum likelihood) point estimate of the dynamics or, with GP residuals, only the residual carries uncertainty. Without full probabilistic inference, the desired accurate uncertainty quantification is missing.

Further, in these works, learning is posed as regression problem, where the input, i.e. the system's kinematic state, is known. On UAVs however, without access to wind-tunnel testing or high fidelity instrumentation, (model-independent) state knowledge cannot be assumed and both the state (a hidden variable) and the model need to be estimated simultaneously. Among different available methods, joint filtering [15], where the model parameters are augmented to the state, allows for the recursive, fully probabilistic treatment of the problem necessary for accurate online uncertainty quantification.

2) Contributions: We leverage semi-parametric modeling and filtering-based, online probabilistic inference to learn the nonlinear dynamics of H-UAVs with both high accuracy and credibility.

- We propose a semi-parametric model structure that is tailored to improve accuracy of (H)-UAV dynamics in critical and difficult-to-model flight phases such as the hover-cruise transitions of tilt-wing aircraft.
- We show how to use the model formulation within a joint unscented Kalman filter (UKF) for credible online, nonlinear state- and model estimation. Particular attention is paid to the inclusion of the selected sparse pseudo-input Gaussian process (SPGP) residual terms in the filter.
- We validate the accuracy and credibility of the proposed approach on the example of learning the highly nonlinear 6DoF dynamics of a tiltwing H-UAV in real-time, using data from both simulated and real systems.

Improvements in accuracy at relevant states and a significant increase in credibility over the residual-free baseline model suggests that the proposed learning framework is a viable candidate to enable future efforts in probabilistically safe control and active exploration for dynamics learning.

II. Problem Statement

We focus on semi-parametric modeling and credible online learning for H-UAVs, which are challenging as they exhibit highly non-linear and only partially observable process dynamics. While the presented inference approach could be applied to other state-space problems, accurate probabilistic inference is known to be worst-case intractable and approximations have to be validated.

For many applications, UAVs are sufficiently modeled as a single rigid body, subject to gravity $\mathbf{g}$ as well as forces $\mathbf{F}$ and moments $\mathbf{M}$ from aerodynamic and inertial/gyroscopic sources (1). In (1), the state $\mathbf{x} \in \mathbb{R}^{n_x}$ is composed of position and ground-speed ($_I\mathbf{r}$, $_B\mathbf{v}$), body attitude and angular rate ($\mathbf{q}_{IB}$ or $\mathbf{R}_{IB}$, $_B\boldsymbol{\omega}$), actuator states $\boldsymbol{\delta}$ and inertial wind $_I\mathbf{v}_w$. $\mathbf{u} \in \mathbb{R}^{n_u}$ is the control input, and the (potentially uncertain) dynamics are parameterized with nominal $\boldsymbol{\Theta} \in \mathbb{R}^{n_\Theta}$ (including mass m and moment of inertia $_B\mathbf{I}$) and residual $\boldsymbol{\Phi} \in \mathbb{R}^{n_\Phi}$ parameters[1].

$$ {}_I\dot{\mathbf{r}} = \mathbf{R}_{IB} \cdot {}_B\mathbf{v} \tag{1a} $$

$$ \dot{\mathbf{q}}_{IB} = \frac{1}{2}\mathbf{q}_{IB} \otimes {}_B\boldsymbol{\omega} \tag{1b} $$

$$ {}_B\dot{\mathbf{v}} = m^{-1}{}_B\mathbf{F}(\mathbf{x};\boldsymbol{\Theta},\boldsymbol{\Phi}) - {}_B\boldsymbol{\omega} \times {}_B\mathbf{v} + {}_B\mathbf{g} \tag{1c} $$

$$ {}_B\dot{\boldsymbol{\omega}} = {}_B\mathbf{I}^{-1}\left({}_B\mathbf{M}(\mathbf{x};\boldsymbol{\Theta},\boldsymbol{\Phi}) - {}_B\boldsymbol{\omega} \times {}_B\mathbf{I}_B\boldsymbol{\omega}\right) \tag{1d} $$

$$ \dot{\boldsymbol{\delta}} = f_\delta(\boldsymbol{\delta}, \mathbf{u}) \tag{1e} $$

$$ {}_I\dot{\mathbf{v}}_w = \mathbf{n}_w \tag{1f} $$

$$ \mathbf{y} = h(\mathbf{x};\boldsymbol{\Theta},\boldsymbol{\Phi}) + \mathbf{n}_y, \tag{1g} $$

$\mathbf{F}$ and $\mathbf{M}$ further include the state-dependent uncertainty carried by the SPGP-based residuals, see III-C. Actuator states $\boldsymbol{\delta}$ (servo positions, propeller rotational speeds) are propagated using separate, e.g., rate-constrained first order systems f_δ, and the wind dynamics are approximated as random walk with $n_{w,i} \sim \mathcal{N}(0, \sigma_{w,i})$. Leading subscript denote body $_B()$ and inertial $_I()$ frames of reference, vector- and matrix quantities are in bold face.

The system is assumed to be observed through noisy measurements $\mathbf{y}$, with noise $n_{y,i} \sim \mathcal{N}(0, \sigma_{y,i})$, from 1) global navigation satellite system (GNSS) position and ground-speed, angular rates and linear acceleration from an inertial measurement unit (IMU) and 3D magnetic field as well as 2) 1D airspeed $v_{a,x}$ and possibly angle-of-attack and sideslip. The former set is central to the estimation of the inertial pose and twist, while the latter provides air-relative quantities relevant for wind estimation.

III. Methodology

To estimate the dynamics of the partially observable process defined above, we follow the practice of joint filtering [15] by augmenting the process model (1) with $[\dot{\boldsymbol{\Theta}}, \dot{\boldsymbol{\Phi}}] = 0$ and treating the model parameters as part of the state $\mathbf{x}' = [\mathbf{x}, \boldsymbol{\Theta}, \boldsymbol{\Phi}]$.

A. Nominal parametric model

The forces and moments in (1), are approximated using common parametric models for the aerodynamic lift- and drag of airfoils and the fuselage as well as propeller forces and torques. For a lack of space, we have to refer to detailed descriptions of these components in previous work [16] and aerodynamic modeling literature, e.g. [4], [17], [18].

Several H-UAVs types, such as the presented tiltwing and some tailsitter designs [19], make use of the forced flow on wing segments located behind propellers. The additional lift, drag and reduced flow-separation effects on these segments are relevant to the dynamics and should not be neglected. A first approximation is obtained by superposition of the

[1]While the approach is conceptually *semi-parametric*, when implemented it is conventional to use a finite parameterization to avoid growth with data. The model, like other sparse GPs and neural networks, is still non-parametric in the sense of being entirely data-driven.

free-stream airflow and the propeller slipstream [16], [19] as well as terms lumped to propeller forces, cf. *ram-drag* in [18]. Due to the complex (turbulent etc.) flow phenomena, however, more accurate yet general modeling is difficult and, instead, data-driven residual terms are included to further improve accuracy.

B. Sparse Gaussian process residuals

Assuming the rigid-body assumption holds, and in the absence of external disturbances, any non-reducible errors in the nominal formulation of (1) can be attributed to either structural errors in $\mathbf{F}$ and $\mathbf{M}$ or the overall reduced order of the dynamics model due to, e.g. the disregard of the state of the fluid surrounding the H-UAV. In this work, we compensate for the former, by adding residual components $\mathbf{r}_F(\mathbf{x};\mathbf{\Phi})$, $\mathbf{r}_M(\mathbf{x};\mathbf{\Phi})$ where deemed necessary.

GP-based residuals allow for flexible function approximation with encoded uncertainty about unseen parts. However, recursive state and model estimation using parameterized dynamics combined with standard GP residual terms is generally not possible. Since GP residuals, model parameters and states all enter the observation model, the GP targets will become correlated with parameters and states through observations. For inference, it would be required to keep track of the increasing set of GP targets and their joint distribution with the state and model parameter estimates.

A more practical solution, which allows for recursive updates and which offers constant update and prediction complexity, is obtained by adopting the SPGP formulation from [20]. It approximates the full inference problem by considering a finite, constant-size set of values $\mathbf{\Phi}$ (*pseudo-targets*) of the GP at selected input locations $\mathbf{x}_{GP} \in \mathbf{X}_{GP}$ (*pseudo-inputs*). The SPGP can thus be seen as a function parameterized by $\mathbf{\Phi}$, with (noise-free) predictions $r \sim \mathcal{N}(\mu_r, \Sigma_r)$ given by the standard GP observation likelihood [21]

$$\begin{aligned}\mu_r(\mathbf{x}^*) &= k_n(\mathbf{x}^*)^T K_n^{-1}\mathbf{\Phi}\\ \Sigma_r(\mathbf{x}^*) &= k(\mathbf{x}^*,\mathbf{x}^*) - k_n(\mathbf{x}^*)^T K_n^{-1} k_n(\mathbf{x}^*)\end{aligned}$$

and $[k_n(\mathbf{x}^*)]_i = k(\mathbf{x}^*, \mathbf{x}_{GP,j})$, $[K_n]_{ij} = k(\mathbf{x}_{GP,i}, \mathbf{x}_{GP,j})$ the Gram matrices with kernel $k(\cdot,\cdot)$ and query location $\mathbf{x}^*$.

C. Filtering with SPGP residuals

Given its parametric form, updating the SPGP ties naturally into the adopted filtering framework where we estimate $\mathbf{\Phi}$ along with other model parameters $\mathbf{\Theta}$ and the system state $\mathbf{x}$ to capture the full joint probability $p(\mathbf{x}, \mathbf{\Theta}, \mathbf{\Phi})$. While the residual's mean μ_r is, technically, just an additional parametrized function, we leverage the covariance Σ_r to retain the key property of a GP to provide an input (i.e. state) dependent uncertainty that grows with increasing distance to observation locations. It will therefore remain uncertain away from $\mathbf{x}_{GP}$, even if $\hat{\mathbf{\Phi}}$ converges. This is desirable, since it reflects the inability to make global residual predictions with information compressed to the sparse finite set $\mathbf{\Phi}$. Formally, this is achieved in the process and measurement model by treating the residual probabilistically as $r \sim \mathcal{N}(\mu_r, \Sigma_r)$ with Σ_r depending on both the belief $[\hat{\mathbf{x}}, \hat{\mathbf{\Phi}}] \sim \mathcal{N}(\mu_{\mathbf{x},\mathbf{\Phi}}, \Sigma_{\mathbf{x},\mathbf{\Phi}})$ in (2b) and the GP prediction uncertainty in (2c)—which persists away from $\mathbf{x}_{GP}$ and represents the uncertainty in the residual where the SPGP cannot learn.

$$\mu_r(\hat{\mathbf{x}}, \hat{\mathbf{\Phi}}) = k_n(\mu_{\mathbf{x}})^T K_n^{-1} \mu_{\mathbf{\Phi}}, \tag{2a}$$

$$\Sigma_r(\hat{\mathbf{x}}, \hat{\mathbf{\Phi}}) = J_{\mu_r}^T \Sigma_{\mathbf{x},\mathbf{\Phi}} J_{\mu_r} \tag{2b}$$

$$+ k(\mu_{\mathbf{x}}, \mu_{\mathbf{x}}) - k_n(\mu_{\mathbf{x}})^T K_n^{-1} k_n(\mu_{\mathbf{x}}) \tag{2c}$$

(2b) depicts a linearized uncertainty mapping with Jacobian $J_{\mu_r} = \partial \mu_r / \partial [\hat{\mathbf{x}}, \hat{\mathbf{\Phi}}]$. This is implicitly accounted for through μ_r when using sampling-based uncertainty propagation, e.g. within a UKF.

D. Performance criteria

We use the mean absolute error (MAE) to measure the average accuracy of predicted linear and angular accelerations, and the inclination index (II) and non-credibility index (NCI) [22] to evaluate the quality of the associated uncertainty predictions. These metrics are obtained from a set of sampled filter realizations j per experiment, see IV-C.

The II and NCI evaluate a 'difference' between the distributions of the actual- and the expected prediction errors. Let $\tilde{x}$ be the model prediction error. The matrix of mean squared prediction errors, P^*, is estimated from M sampled model predictions as $P^* = \frac{1}{M}\Sigma_{j=1}^{M} \tilde{x}_j \tilde{x}_j'$. It is compared with the predicted variance P by means of the *credibility ratio* $\rho_j = \tilde{x}_j^T P_j^{-1} \tilde{x}_j / \tilde{x}_j^T P^{*-1} \tilde{x}_j$ [22]. If the predicted distribution of estimation errors is 'wider' than the true/sampled one, the *credibility ratio* will be less than one (the model is 'pessimistic'). This is not accurate, but still safe for bounding the estimated system dynamics. A 'narrower' (optimistic) estimate would also not be accurate, but could additionally be unsafe as it is incorrectly over-confident. $II(k) = \frac{10}{M}\Sigma_{j=1}^{M} \log_{10}(\rho_j)$ and $NCI(k) = \frac{10}{M}\Sigma_{j=1}^{M} |\log_{10}(\rho_j)|$ reflect this optimism or pessimism and overall accuracy by considering the sum of signed or absolute logs, respectively. For high credibility, the goal is to minimize NCI (≥ 0) while, for safety, a tendency towards conservative uncertainty estimates ($II \leq 0$) is desired.

IV. Experimental Validation

We demonstrate real-time inference with the proposed framework for learning the full 6-DoF nonlinear dynamics of a small tiltwing H-UAV from on-board measurements. In previous work [6], we showed that a batch-optimized (nominal) parametric model provided a good approximation of the translational dynamics. However, it failed at predicting the pitch dynamics during transitions. Here, we investigate possible improvements attained through semi-parametric learning, both in terms of accuracy and model credibility through probabilistic inference.

A. Tiltwing Model

The considered tiltwing H-UAV (Fig. 1) is a custom platform [23] with a wingspan of $2\mathrm{m}$, a take-off weight of $7\mathrm{kg}$, and a cruise speed of up to $25\mathrm{m/s}$.

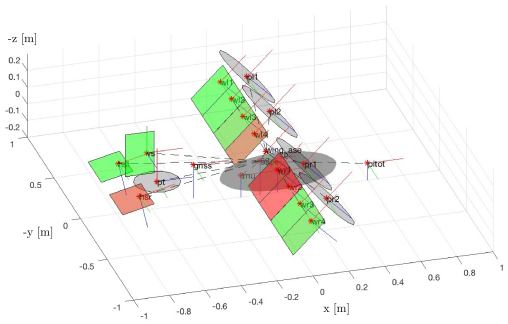

Fig. 2. Illustration of the component build-up underlying the dynamics model. Individual airfoil segments, propellers, and the fuselage contribute to the net wrench based on local inflow conditions and, if applicable, actuator states.

1) Dynamics: The process model of the filter is based on (1) with aerodynamic forces and moments summed from individual contributions of airfoil segments, propellers and the fuselage as detailed in III-A and [16], see Fig. 2. Notably, we apply the described propeller-wing interaction modeling for the fully blown wing. Servo and propeller dynamics are modeled as rate- and input-constrained first order systems. Overall, $n_x = 26$ including the states of all $n_u = 11$ actuators.

A set of $n_\Theta = 27$ nominal model parameters is estimated in the following experiments. It includes the moments of inertia $(I_{xx}, I_{yy}, I_{zz}, I_{xz})$, center of gravity location $(-d_x, -d_z)$, time constants for control surfaces, main- and tail-propeller $(\tau_{cs,ele}, \tau_{cs,ail}, \tau_{pm}, \tau_{pt})$, propeller thrust coefficient $c_{T,J}$, aerodynamic coefficients for the wing $(c_{L,0}, c_{L,\alpha}, c_{L,\delta_{ail}}, c_{L,fp}, c_{D,0}, c_{D,\alpha^2}, c_{D,fp})$ and a subset thereof for the horizontal and vertical stabilizers. Finally, we include correction factors for ram-drag (κ_{ram}), propeller slipstream strength (κ_{slip}) and fuselage drag $(\kappa_{cd,x}, \kappa_{cd,y}, \kappa_{cd,z})$. Further details can be found in [6], [16].

SPGP residual terms are added to less well-fitted components of the parametric model, i.e., the body-z force Z as well as the roll, pitch, and yaw moments L, M, N using inputs known to affect these respective components. The overall number of pseudo-inputs was maximized within computational constraints, see IV-B.

$r_Z(v_{a,x}, \delta_{pm}, \delta_w; \mathbf{\Phi}_Z)$	$n_{\Phi_Z} = 20$
$r_L(\Delta\delta_{pm}, \Delta\delta_{ail}, \delta_w; \mathbf{\Phi}_L)$	$n_{\Phi_L} = 10$
$r_M(v_{a,x}, \delta_{pm}, \delta_w; \mathbf{\Phi}_M)$	$n_{\Phi_M} = 20$
$r_N(v_{a,x}, v_{a,y}, \omega_z; \mathbf{\Phi}_N)$	$n_{\Phi_N} = 10$

The residual's pseudo input locations are uniformly distributed across the projection of the flight envelope onto the corresponding inputs, i.e., they are placed close to states the UAV typically operates at. These operating points can be obtained from existing flight data or a model prior. We use squared-exponential kernels with a diagonal covariance. Their length scales and variance are not specifically optimized. As a preliminary heuristic and for the purpose of demonstrating the overall approach, we choose the length scales proportionally to the distance between pseudo-input points. If they are chosen too small, the SPGP can only learn close to the $\mathbf{x}_{GP}$, if they are too large, there is a risk for overly confident extrapolation.

Overall, $n_{x'} = 113$ for the semi-parametric model M_{sp} and $n_{x'} = 53$ for the nominal parametric model M_p. The dynamics used to simulate the platform and generate synthetic data feature, for most components, an identical structure to the filter process model. Differences stem from simplifications made to the process model, and purposely introduced mismatches in the simulated propeller-wing interaction to be captured by the residual terms of the semi-parametric model.

2) Measurements: Obtained from the onboard (simulated and real) Pixhawk 4 autopilot [24] with its IMU, barometer, and magnetometer, as well as a GNSS unit and a 1-D airspeed probe. For convenience and lack of instrumentation, respectively, magnetic field $\mathbf{B}$ and angle-of-attack AoA are reconstructed from the flight controller's internal kinematic state estimate.

signal	sensor	f_{upd}	σ
${}_I\mathbf{r}$	GNSS	5Hz	$[1, 1, 1]$ m
${}_I\mathbf{v}$	(ublox Neo-M8N)	5Hz	$[.2, .2, .2]$ m/s
${}_B\boldsymbol{\omega}$	IMU	20Hz	$[.025, .025, .025]$ rad/s
${}_B\mathbf{a}_{IMU}$	(BMI055/ICM-20689)	20Hz	$[.25, .37, .53]$ m/s^2
${}_B\mathbf{B}$	pseudo-measurement	20Hz	$[.1, .1, .1]$ $[-]$
v_a	Pitot, (flow-based, SDP33)	5Hz	1 m/s
AoA	pseudo-measurement	5Hz	.2 rad

B. Filter

In all experiments, filtering is performed using a UKF due to its superior performance on nonlinear problems compared to, e.g., an extended Kalman filter (EKF), which was found to diverge for a majority of initial conditions with the considered dynamics. We use identical settings for process and measurement noises for comparability amongst all experiments. The C++ implementation of the UKF with a maximum of $n_{\hat{x}'} = 113$ states for the semi-parametric model achieves an update rate of $>$ 20Hz on a laptop-grade CPU (2.9GHz Intel-i7, single core). Hence, real-time inference is possible with the employed time-discretization of 0.05s in all experiments. Note that the complexity of a UKF scales with $\mathcal{O}(L^2)$ in number L of parameters/pseudo-inputs [25].

C. Evaluation methodology

The baseline parametric M_p and the semi-parametric model M_{sp} are estimated using recorded input and measurement trajectories $\mathbf{u}(k)$, $\mathbf{y}(k)$ from simulated and real flights. At any time k during the filtering process on a given trajectory T, the instantaneous belief of model $[\hat{\mathbf{\Theta}}, (\hat{\mathbf{\Phi}})]_{k|k}$ can be evaluated based on its accuracy and credibility in predicting linear $(\mathbf{a})$ and angular accelerations $(\boldsymbol{\alpha})$

$$ {}_B\hat{\mathbf{a}}_{IMU} = \hat{m}^{-1}\,{}_B\mathbf{F}(\mathbf{x}; \hat{\mathbf{\Theta}}, \hat{\mathbf{\Phi}}) + {}_B\hat{\mathbf{a}}_c \tag{3a} $$

$$ {}_B\hat{\boldsymbol{\alpha}} = {}_B\hat{\mathbf{I}}^{-1}\left({}_B\mathbf{M}(\mathbf{x}; \hat{\mathbf{\Theta}}, \hat{\mathbf{\Phi}}) - {}_B\boldsymbol{\omega} \times {}_B\hat{\mathbf{I}}\,{}_B\boldsymbol{\omega}\right) \tag{3b} $$

at some state $\mathbf{x}$. ${}_B\hat{\mathbf{a}}_c = {}_B\hat{\dot{\boldsymbol{\omega}}} \times {}_B\hat{\mathbf{d}} + {}_B\boldsymbol{\omega} \times ({}_B\boldsymbol{\omega} \times {}_B\hat{\mathbf{d}})$ are the Coriolis and centrifugal accelerations measured by an IMU offset from the center of gravity by $\mathbf{d}$. The prediction errors are formed with the measured acceleration

${}_B\mathbf{a}_{IMU}(\mathbf{x})$ and the time derivative of the measured angular rate ${}_B\dot{\boldsymbol{\omega}}_{IMU}(\mathbf{x})$. To evaluate II and NCI, the required error statistics are approximated by sampling 100 prior model means $\boldsymbol{\mu}^{j}_{\Theta,\Phi,0|0} \sim \mathcal{N}(\boldsymbol{\mu}^{*}_{\Theta,\Phi}, \Sigma_{\Theta,\Phi,0|0})$, and therefore filtering runs j, per learning trajectory. Assuming $\boldsymbol{\mu}^{*}_{\Theta,\Phi}$ lies close to the ground truth, we therefore initialize each filter j with a credible estimate.

In (3), when evaluating the models learned from real flight data, we do not have access to the ground-truth state $\mathbf{x}$. Instead we need to rely on an estimate $\hat{\mathbf{x}}$ which will, since its obtained jointly with the model, inevitably be biased towards the model to a degree depending on the filter tuning. E.g., the joint estimate of the wind and the dynamics might give correct (instantaneous) acceleration predictions, while both wind and dynamics model are, however, wrong. It is then specific to the application, whether an accurate motion prediction or an accurate model itself is desired. We therefore present two choices for obtaining $\hat{\mathbf{x}}$ when evaluating models learned from real data. The first uses, for each model structure $m \in \{M_p, M_{sp}\}$ and filter realization j, the individual, UKF-smoothed [26] state estimate $\hat{\mathbf{x}}^{m,j}_k|\mathbf{y}_{0:N}$. Alternatively, to counteract potential biases introduced with this method, we also evaluate the model using the sample mean and covariance of the smoothed state means $\boldsymbol{\mu}^{m,j}_{x,k}|\mathbf{y}_{0:N}$ over all model structures and filter realizations. To reduce potential biases of $\hat{\mathbf{x}}$ in the first place, the filter noise values are tuned to obtain similar smoothed state trajectories between models and filter realizations.

For the actual choice of $\mathbf{x}$ in (3), we consider the full flown *training* trajectory T_{train} and, in simulation, an additional *test* trajectory T_{test}. Hence, when evaluating an intermediate belief of the model, both past/seen, as well as future, possible unseen, states are included. This allows to simultaneously investigate two desired properties of the learning method: 1) to not forget about the past (and locally overfit) as well as 2) to reliably extrapolate to new states.

D. Results

For ease of comparison and presentation, we summarize the statistics of selected performance metrics over all states of a given trajectory using boxplots, cf. Fig. 3 and Fig. 4. The MAE, as well as the II and NCI are evaluated per state, and are themselves statistics (averages) over all 100 filter realizations. For the predicted Mahalanobis distances d_M (the numerator in ρ, see III-D), the boxplots summarize over both states and filter realizations j.

1) Learning from simulated data: In simulation, T_{train} is (heuristically) designed to achieve observability of approximately steady translational- and decoupled angular dynamics. The H-UAV gradually transitions from hover- to cruise flight, and back. At several intermediate wing tilt angles, all three axes roll, pitch and yaw are individually excited using step-like inputs. T_{test} on the other hand, is more dynamic, involving rapid transitions and coupled changes in the attitude references.

Fig. 3 shows the accuracy and credibility attained with the filtered $()^f$ beliefs of the nominal parametric (M^f_p) and

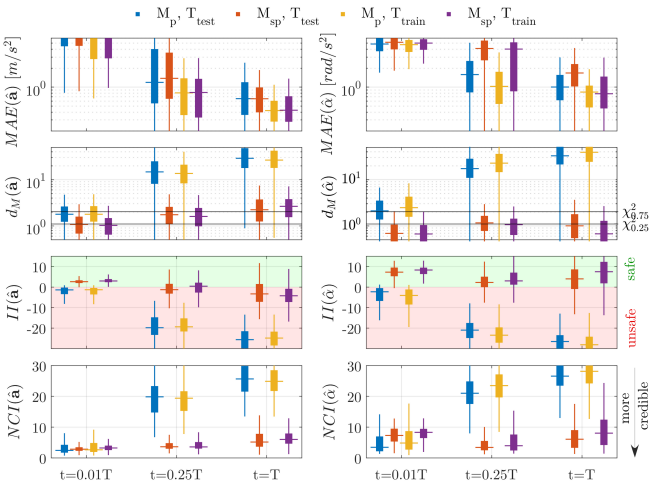

Fig. 3. Initial, intermediate, and final prediction error and credibility statistics for the models learned from simulated data.

the semi-parametric (M^f_{sp}) models at three different times during learning. Medians of axis-wise MAEs and results for batch-optimized $()^b$ models M^b_p, M^b_{sp} are detailed in Table I.

At $t = 0.25T$, i.e. a quarter into the training trajectory, errors are still comparably high, since only a subset of states has been visited. E.g., the dynamics of cruise-flight, and relevant model parameters such as lift coefficients, have not, or only partially become observable at this point. Compared to M^f_{sp}, however, the credibility of M^f_p is already significantly degraded and will reduce further as model beliefs converge with additional data until $t = T$, rendering M^f_p overconfident and therefore unsafe. M^f_{sp} on the other hand, despite a slower learning rate and similar or lower accuracy on the angular dynamics, retains its comparably high credibility. The lower accuracy of M^f_{sp} on the test set might indicate over-fitting to the training set, albeit in a credible way, i.e., where larger prediction errors are paired with higher prediction uncertainty. Since the simulated dynamics and M_p are still very similar, the additional expressiveness of M_{sp} is likely to hamper extrapolation accuracy, with residuals making up for parts that M_p could, and ideally should, explain. A similar tendency is observed when comparing the absolute error statistics of the corresponding batch-optimized models M^b_p, M^b_{sp} in Table I.

2) Learning from real data: As described in IV-C, when using real data, the hidden state $\mathbf{x}$ needs to be estimated. In the following, $()^{ego}$ refers to the use of an estimate $\hat{\mathbf{x}}$ individual to each filter realization, and $()^{uni}$ to the use of

TABLE I

MEDIAN PREDICTION MAE AT $t = T$ ON SIMULATED DYNAMICS

trajectory	model	$\mathbf{a}$ $[m/s^2]$				$\boldsymbol{\alpha}$ $[rad/s^2]$			
		a_x	a_y	a_z	$\mathbf{a}$	α_x	α_y	α_z	$\boldsymbol{\alpha}$
T_{train}	M^f_p	0.35	0.01	**0.18**	**0.42**	**0.04**	0.69	**0.04**	0.83
	M^f_{sp}	**0.22**	0.01	0.31	0.43	0.07	**0.65**	0.06	**0.78**
T_{test}	M^f_p	0.29	**0.03**	**0.45**	0.65	**0.15**	**0.79**	**0.13**	**1.00**
	M^f_{sp}	**0.19**	0.04	0.54	0.65	0.18	1.33	0.23	1.68
	M^b_p	0.21	0.03	0.29	0.43	0.12	**0.38**	0.09	**0.56**
	M^b_{sp}	**0.18**	**0.02**	**0.24**	**0.39**	0.12	0.66	0.09	0.83

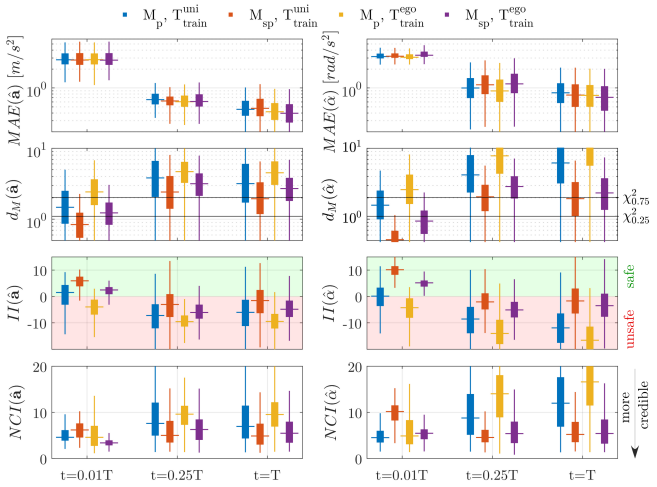

Fig. 4. Initial, intermediate, and final prediction error and credibility statistics for the models learned from real data.

TABLE II

MEDIAN PREDICTION MAE AT $t = T$ ON REAL DYNAMICS

trajectory	model	$\mathbf{a}$ $[m/s^2]$				$\boldsymbol{\alpha}$ $[rad/s^2]$			
		a_x	a_y	a_z	$\mathbf{a}$	α_x	α_y	α_z	$\boldsymbol{\alpha}$
T^{ego}_{train}	M^f_p	0.17	0.15	0.24	0.42	**0.28**	0.34	0.36	0.76
	M^f_{sp}	**0.16**	**0.13**	**0.23**	**0.40**	0.29	**0.30**	**0.31**	**0.71**
T^{uni}_{train}	M^f_p	0.17	**0.14**	**0.29**	**0.46**	0.34	0.41	0.36	0.84
	M^f_{sp}	0.17	0.15	0.31	0.47	**0.31**	**0.36**	**0.32**	**0.77**
	M^b_p	0.11	0.14	0.23	0.37	0.30	0.32	0.33	0.73
	M^b_{sp}	**0.09**	0.14	**0.18**	**0.32**	**0.25**	**0.27**	**0.30**	**0.64**

an averaged, single $\hat{\mathbf{x}}$. The trajectory T_{train} involves two successive, rapid transitions with mostly smooth changes in attitude references and hence covers states common to a typical mission profile.

Fig. 4 shows again an overall similar final MAE statistic between M^f_p and M^f_{sp} and a significantly higher credibility attained by M^f_{sp}. On T^{ego}_{train}, M^f_{sp} achieves lower mean and median MAE for both linear- and angular accelerations, see Table II. While the average difference over the full trajectory is small, a closer inspection of particular phases of T_{train} reveals consistent model improvements attained with M^f_{sp}. Notably, we observe the desired gain in accuracy for the pitch dynamics during transition, where we added a residual specifically to compensate for the deficiencies of the nominal model found in previous work [6]. In Fig. 5 we compare predictions on states of a transition back from cruise to hover. This part of the flight requires considerable pitch stabilization effort, due to strong deceleration forces coupled with shifts in the center of pressure. Accurate modeling is thus particularly relevant to enable safe control and, e.g., prevent control saturation by reducing transition speeds. Both model realizations $M^{f_j}_p$, $M^{f_j}_{sp}$ fail to predict α_y at unseen future states initially, but M^f_{sp} does so, on average, with higher credibility. Once the transition is fully observed (middle), the errors are reduced. Compared to M^f_{sp}, however, all realizations of M^f_p consistently overestimate α_y shortly after $t = 530$s, as shown by the distributions of all 100 filter's mean predictions on T^{uni}_{train} in Fig. 5, bottom.

Overall, M^f_{sp} mainly improves accuracy on the angular dynamics, see Table II. This agrees with previous experience [6], which indicates sufficient descriptiveness of a similar M_p for the translational dynamics of a tiltwing. Furthermore, by the placement of residuals, M_{sp}'s expressiveness and potential for improvement over M_p, is mainly increased on the angular dynamics in the first place.

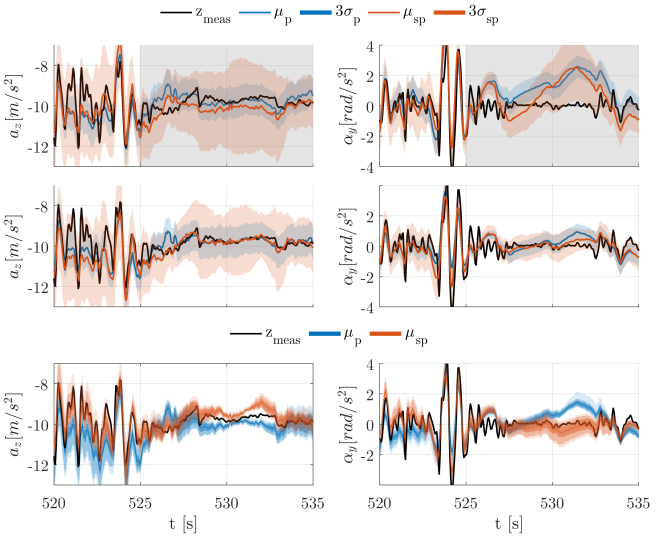

Fig. 5. Vertical linear (a_z) and pitch angular acceleration (α_y) predictions of single realizations $M^{f_j}_p$, $M^{f_j}_{sp}$ before ($t = 525$s, top row) and after ($t = 550$s, middle row) observing the transition dynamics from cruise- to hover flight. The grey shaded part delimits future, unseen states. Bottom row: 68th-95th-99th percentiles of predicted means over all 100 filter realizations at $t = 550$s, using T^{uni}_{train}, showing that M_{sp} performs better as desired.

V. CONCLUSIONS AND FUTURE WORK

We proposed a semi-parametric framework for estimation of H-UAV dynamics that allows us to add SPGP residuals for critical and difficult-to-model flight phases such as the hover-cruise transitions of tiltwing H-UAV. We showed that the SPGP residual terms allow the model to maintain higher credibility throughout the learning process, i.e. with both low- and high dynamics observability, while the corresponding residual-free model quickly becomes over-confident and therefore unsafe. We also showed that the added expressiveness of residual components leads to improved accuracy as intended in the difficult-to-model angular dynamics of a real H-UAV. Finally, we showed that our UKF-based filter implementation allows online, real-time learning at > 20Hz from real flight data on a laptop-grade CPU. Future steps include evaluation of the learning framework for use with safe predictive control, as well as further consideration of the number and placement of the SPGP psuedo-input points.

ACKNOWLEDGMENTS

This work was funded in part by the Amazon Research Awards program, Dufour Aerospace and a Wallenberg Foundation and WASP Postdoctoral Scholarship.

REFERENCES

[1] T. Koller, F. Berkenkamp, M. Turchetta, J. Boedecker, and A. Krause, "Learning-based Model Predictive Control for Safe Exploration and Reinforcement Learning," Tech. Rep.

[2] S. Patil, G. Kahn, M. Laskey, J. Schulman, K. Goldberg, and P. Abbeel, "Scaling up Gaussian Belief Space Planning through Covariance-Free Trajectory Optimization and Automatic Differentiation," Tech. Rep.

[3] J. Van Den Berg, S. Patil, and R. Alterovitz, "Motion planning under uncertainty using iterative local optimization in belief space," *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1263–1278, 2012.

[4] M. S. Selig, "Modeling Full-Envelope Aerodynamics of Small UAVs in Realtime," in *AIAA Atmospheric Flight Mechanics 2010 Conference*, Aug.

[5] C. Olsson, S. L. Verling, T. Stastny, and R. Siegwart, "Full Envelope System Identification of a VTOL Tailsitter UAV," in *AIAA Scitech 2021 Forum*. Reston, Virginia: American Institute of Aeronautics and Astronautics, jan 2021. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2021-1054

[6] D. Rohr, M. Studiger, T. Stastny, N. R. J. Lawrance, and R. Siegwart, "Nonlinear model predictive velocity control of a vtol tiltwing uav," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5776–5783, 2021.

[7] S. Riedel and F. Stulp, "Comparing semi-parametric model learning algorithms for dynamic model estimation in robotics," *arXiv preprint arXiv:1906.11909*, 2019.

[8] I. Georgiev, C. Chatzikomis, T. Völkl, J. Smith, and M. Mistry, "Iterative semi-parametric dynamics model learning for autonomous racing," *arXiv preprint arXiv:2011.08750*, 2020.

[9] J. Smith and M. Mistry, "Online simultaneous semi-parametric dynamics model learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2039–2046, 2020.

[10] A. Gijsberts and G. Metta, "Real-time model learning using incremental sparse spectrum gaussian process regression," *Neural Networks*, vol. 41, pp. 59–69, 2013, special Issue on Autonomous Learning.

[11] D. Romeres, M. Zorzi, R. Camoriano, and A. Chiuso, "Online semi-parametric learning for inverse dynamics modeling," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 2945–2950.

[12] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious model predictive control using gaussian process regression," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2736–2743, 2020.

[13] W. Zhang, M. Tognon, L. Ott, R. Siegwart, and J. Nieto, "Active model learning using informative trajectories for improved closed-loop control on real robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4467–4473.

[14] G. Torrente, E. Kaufmann, P. Föhn, and D. Scaramuzza, "Data-driven mpc for quadrotors," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3769–3776, 2021.

[15] E. Wan, R. Van Der Merwe, and A. Nelson, "Dual estimation and the unscented transformation," *Advances in neural information processing systems*, vol. 12, 1999.

[16] D. Rohr, T. Stastny, S. Verling, and R. Siegwart, "Attitude and Cruise Control of a VTOL Tiltwing UAV," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, jul 2019.

[17] B. Stevens and F. Lewis, *Aircraft Control and Simulation*. Wiley, 2003. [Online]. Available: https://books.google.ch/books?id=T0Ux6av4btIC

[18] M. S. Selig, "Modeling Propeller Aerodynamics and Slipstream Effects on Small UAVs in Realtime," in *AIAA Atmospheric Flight Mechanics 2010 Conference*, Aug., pp. 1–23.

[19] S. Verling, B. Weibel, M. Boosfeld, K. Alexis, M. Burri, and R. Siegwart, "Full Attitude Control of a VTOL tailsitter UAV," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May, pp. 3006–3012.

[20] E. Snelson and Z. Ghahramani, "Sparse gaussian processes using pseudo-inputs," in *Advances in Neural Information Processing Systems*, Y. Weiss, B. Schölkopf, and J. Platt, Eds., vol. 18. MIT Press, 2005.

[21] C. E. Rasmussen, *Gaussian Processes in Machine Learning*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 63–71.

[22] X. R. Li and Z. Zhao, "Measuring estimator's credibility: Noncredibility index," in *2006 9th International Conference on Information Fusion*, 2006.

[23] (2022, Sept.) Dufour Aerospace. [Online]. Available: https://www.dufour.aero/

[24] (2019, Jan.) Pixhawk Autopilot Research Project. [Online]. Available: https://pixhawk.org/

[25] R. Van der Merwe and E. Wan, "The square-root unscented kalman filter for state and parameter-estimation," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221)*, vol. 6, 2001, pp. 3461–3464 vol.6.

[26] S. Särkkä, "Unscented rauch–tung–striebel smoother," *IEEE Transactions on Automatic Control*, vol. 53, no. 3, pp. 845–849, 2008.