

Set-Based State Estimation With Probabilistic Consistency Guarantee Under Epistemic Uncertainty

Shen Li , Theodoros Stouraitis , Michael Gienger , *Member, IEEE*, Sethu Vijayakumar , and Julie A. Shah 

Abstract—Consistent state estimation is challenging, especially under the epistemic uncertainties arising from learned (nonlinear) dynamic and observation models. In this work, we propose a set-based estimation algorithm, named Gaussian Process-Zonotopic Kalman Filter (GP-ZKF), that produces zonotopic state estimates while respecting both the epistemic uncertainties in the learned models and aleatoric uncertainties. Our method guarantees probabilistic consistency, in the sense that the true states are bounded by sets (zonotopes) across all time steps, with high probability. We formally relate GP-ZKF with the corresponding stochastic approach, GP-EKF, in the case of learned (nonlinear) models. In particular, when linearization errors and aleatoric uncertainties are omitted and epistemic uncertainties are simplified, GP-ZKF reduces to GP-EKF. We empirically demonstrate our method’s efficacy in both a simulated pendulum domain and a real-world robot-assisted dressing domain, where GP-ZKF produced more consistent and less conservative set-based estimates than all baseline stochastic methods.

Index Terms—State estimation, nonlinear filtering, robust/adaptive control, physical human-robot interaction.

I. INTRODUCTION

STATE estimation is critical for robot decision making, especially during human-robot interactive tasks, where physical states can be partially observed due to occlusions [1] and latent mental states can influence the interaction [2], [3]. For a robot to safely interact with a human, such states are typically estimated based on models [4]. Due to the complexity in model specification, human behavior and observation models are usually learned from data [5]. These learned models typically include errors, and any process has stochastic noise; the former is described as

Manuscript received October 6, 2021; accepted February 7, 2022. Date of publication March 1, 2022; date of current version April 12, 2022. This letter was recommended for publication by Associate Editor M. Li and Editor J. Kober upon evaluation of the reviewers’ comments. This work was supported in part by the USA Office of Naval Research under Award ID N00014-18-1-2815, in part by Alan Turing Institute, and in part by Honda Research Institute Europe. (*Shen Li and Theodoros Stouraitis contributed equally to this work.*) (*Corresponding author: Shen Li.*)

Shen Li and Julie A. Shah are with the CSAIL, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: shenli@mit.edu; julie_a_shah@csail.mit.edu).

Theodoros Stouraitis is with the CSAIL, Massachusetts Institute of Technology, Cambridge, MA 02139 USA, and also with the School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, U.K. (e-mail: theostou@mit.edu).

Sethu Vijayakumar is with the School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, U.K., and also with the Alan Turing Institute, London NW1 2DB, U.K. (e-mail: sethu.vijayakumar@ed.ac.uk).

Michael Gienger is with the Honda Research Institute Europe, 63073 Offenbach, Germany (e-mail: michael.gienger@honda-ri.de).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2022.3154802>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2022.3154802




Fig. 1. In a robot-assisted dressing scenario, we deployed our set-based estimator—GP-ZKF—to estimate the visually occluded human elbow position [1]. With human dynamic and observation models learned via Gaussian Processes, GP-ZKF constructs zonotopic state estimates (illustrated with the green box) based on the force measurements at the robot end effector. By handling the epistemic uncertainties in the learned models, GP-ZKF guarantees probabilistic consistency—*i.e.*, the true human elbow positions are bounded by the zonotopes across all time steps, with high probability.

“epistemic uncertainty,” and the latter as “aleatoric uncertainty,” both of which present critical challenges for estimation algorithms. For example, in a robot-assisted dressing scenario (see Fig. 1), the epistemic uncertainty could blind the robot, resulting in overconfidence in an erroneous human state and rendering the robot’s behavior “aggressive” and unsafe. To address this challenge, we developed a set-based estimation algorithm that is able to *conservatively* respect these uncertainties.

We focus on the problem of consistency in state estimation from the view of epistemic uncertainty introduced by learned models. In prior literature, consistency has been analyzed within two paradigms: stochastic paradigm and set-based paradigm. For stochastic methods, such as the Extended Kalman Filter (EKF), a consistent estimate is defined as an unbiased point estimate together with a covariance matching the actual estimation error [6]. In the field of SLAM, the inconsistency issue of EKF-based approaches, such as GP-EKF [7], has been broadly investigated from the views of linearization errors [8], [9] and state unobservability [10], [11]. On the other hand, set-based methods construct sets, as state estimates, where consistency can be interpreted as the sets containing the true states [12]—*i.e.*, bounded estimation error. For known models, where *epistemic* uncertainty can be omitted, set-based methods can produce estimates with guaranteed consistency under aleatoric uncertainty [12], [13]. In contrast, we consider settings in which nonlinear models are learned; in such circumstances, uncertainty

arises from both the learning errors (*epistemic*) and noises (*aleatoric*).

We developed Gaussian Process-Zonotopic Kalman Filter (GP-ZKF), a set-based estimation algorithm with a probabilistic consistency guarantee under learned (dynamic and observation) models. GP-ZKF learns both nonlinear models via Gaussian Processes (GP) and leverages GP's confidence intervals [14] to calibrate the epistemic uncertainties in both models. This extends prior work [15] that assumed bounded epistemic uncertainties in the linear parameter-varying enclosures of the nonlinear models. Similar to set-based estimators [12], [13], but specifically for scenarios with learned models, our approach recursively produces set-based estimates that are represented as zonotopes (a special type of polytope). These zonotopes are designed to respect both epistemic and aleatoric uncertainties and guaranteed to contain the true states across all times steps, with high probability, rendering GP-ZKF consistent when both nonlinear models are learned.

We formally relate our *set-based* approach—GP-ZKF—with the corresponding *stochastic* approach—GP-EKF [7]—and prove that GP-ZKF reduces to GP-EKF if GP-ZKF omits linearization errors and aleatoric, and simplifies epistemic uncertainties. This theoretical connection under nonlinear and learned models extends prior work [16] in relating set-based with stochastic paradigms under linear and known models.

In this paper, we make the following contributions:

- 1) We propose GP-ZKF, a set-based state estimator with a probabilistic consistency guarantee under epistemic and aleatoric uncertainties, in the case where both dynamic and observation models are nonlinear and learned.
- 2) We formally relate GP-ZKF with its stochastic counterpart, GP-EKF [7], under nonlinear and learned models.

We evaluated GP-ZKF in a simulated pendulum domain and a real-world robot-assisted dressing domain; our results show that GP-ZKF provides not only more consistent, but less conservative set-based estimates than the stochastic baselines (GP-EKF, GP-UKF, and GP-PF [7]). To the best of our knowledge, ours is the first work with a probabilistic consistency guarantee for state estimation with learned models.

In Section II, we provide the background, in Section III we present the system formulation, and in Section IV we formally define probabilistic consistency—our focal point for this work. We present our method in Section V, prove its two main theorems in Section VI, and empirically evaluate it in Section VII.

II. BACKGROUND

Here, we introduce our nomenclature and briefly provide the background on intervals, zonotopes, Gaussian Process (GP), and a high-probability bound for Gaussian noise.

Nomenclature: Uppercase symbols denote sets, bold upper case symbols denote matrices, lowercase symbols denote scalars, and bold lowercase symbols denote vectors. Subscripts denote the time and dimension: *e.g.*, $\mathbf{x}_{t,j}$, for example, is the j th dimension of a vector, \mathbf{x} , at time t . Superscripts denote the function a variable is related to—*e.g.*, μ^g is the mean of the function $g(\cdot)$. The Minkowski sum, \oplus , between two sets, X and Y , is defined as $X \oplus Y := \{\mathbf{x} + \mathbf{y} : \mathbf{x} \in X, \mathbf{y} \in Y\}$.

A. Intervals and Zonotopes

An interval $[a, b]$ is defined as $\{x : a \leq x \leq b\}$. A box $\subset \mathbb{R}^n$ is an interval vector, $([a_1, b_1], \dots, [a_n, b_n])^T$. A zero-centered box with radius $\mathbf{r} \in \mathbb{R}^n$ is defined as $[\mathbf{0} \pm \mathbf{r}] \subset \mathbb{R}^n$.

Zonotopes are convex polytopes that are centrally symmetric [17]. A set $Z \in \mathbb{R}^n$ is a zonotope if there exists $(\mathbf{G}_Z, \mathbf{c}_Z) \in \mathbb{R}^{n \times n_\xi} \times \mathbb{R}^n$, such that $Z = \{\mathbf{c}_Z + \mathbf{G}_Z \xi : \xi \in \mathbb{R}^{n_\xi}, \|\xi\|_\infty \leq 1\}$. Here, \mathbf{c}_Z is the center, each column of \mathbf{G}_Z is a generator, and ξ contains all generator variables. In this work, we let $(Z)_c$ refer to the center, \mathbf{c}_Z , and let $(Z)_G$ refer to the generators, \mathbf{G}_Z . We compactly denote $Z = \{\mathbf{G}_Z, \mathbf{c}_Z\}$. Zonotopes are closed under affine transformations and Minkowski sums, both of which can be computed exactly. Formally, $\mathbf{b} \oplus \mathbf{A}Z = \{\mathbf{A}\mathbf{G}_Z, \mathbf{b} + \mathbf{A}\mathbf{c}_Z\}$ and $Z_1 \oplus Z_2 = \{[\mathbf{G}_{Z_1} \ \mathbf{G}_{Z_2}], \mathbf{c}_{Z_1} + \mathbf{c}_{Z_2}\}$. For details about intervals and zonotopes, please refer to [12], [13], [18].

B. Gaussian Process and Confidence Intervals

In this work, unknown functions are learned via GPs. First, for any unknown function $\psi(\mathbf{x}) : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_y}$, we assume that along all dimensions, $j \in J$ with $J = \{1, \dots, n_y\}$, the function outputs, $\psi_j(\cdot)$, are independent of each other. Then, we equivalently reformulate the multi-output function, $\psi(\cdot)$, using a single-output *surrogate* function, $\psi'(\cdot)$ [14]. We define $\psi'(\cdot) : \mathbb{R}^{n_x} \times J \rightarrow \mathbb{R}$, where for each dimension j , we have $\psi'(\mathbf{x}, j) = \psi_j(\mathbf{x})$; this allows us to apply the standard definition of GP with a scalar output and formulate confidence intervals.

We use a GP, denoted by $GP(m^\psi, k^\psi)$, to learn $\psi'(\cdot)$, where the prior mean function, $\mu^\psi(\cdot)$, is set to 0. We assume that the given training data points are corrupted by *i.i.d.* Gaussian noise $\mathbf{v} \in \mathbb{R}^{n_y}$. Formally, for each dimension j , the noise $v_j \sim \mathcal{N}(0, \lambda_v^2)$, where $\lambda_v \in \mathbb{R}$. By conditioning the GP on n data points, we obtain that, for each j , a posterior mean function, $\mu_{n,j}^\psi(\cdot)$, and a posterior variance function, $(\sigma_{n,j}^\psi)^2(\cdot)$.

This work assumes that the true function, $\psi'(\cdot)$, belongs to the reproducing kernel Hilbert space (RKHS) associated with the kernel, k^ψ . The smoothness of $\psi'(\cdot)$ can be measured via its RKHS norm, denoted by $\|\psi'\|_{k^\psi}$ [14]; then, the confidence intervals for the function outputs of $\psi'(\cdot)$ —or, equivalently, of $\psi(\cdot)$ —can be formulated as follows:

Lemma 1 (Lemma 1 in [19]): Let $\delta^\psi \in (0, 1)$. Fix ψ in RKHS with $\|\psi'\|_{k^\psi} \leq B^\psi$. Assume that each dimension of the data of $\psi'(\cdot)$ is corrupted by *i.i.d.* Gaussian noise $\mathbf{v}_j \sim \mathcal{N}(0, \lambda_v^2)$. Let $\beta_n^\psi = B^\psi + \sqrt{2(\gamma_{n \cdot n_y} + \log(1/\delta^\psi))}$. Then, with probability at least $(1 - \delta^\psi)$, jointly for all $n \in \mathbb{N}$, $j = 1, \dots, n_y$, $\mathbf{x} \in \mathbb{R}^{n_x}$, we have that $|\mu_{n,j}^\psi(\mathbf{x}) - \psi_j(\mathbf{x})| \leq \beta_n^\psi \cdot \sigma_{n,j}^\psi(\mathbf{x})$.

This lemma states that with high probability, jointly for each dimension j , the function output of $\psi_j(\cdot)$ is bounded by a confidence interval centered at the posterior mean prediction $\mu_{n,j}^\psi$. The scaling, β_n^ψ , depends upon $\gamma_{n \cdot n_y}$, the information capacity with $n \cdot n_y$ data points for $\psi'(\cdot)$, which can be bounded when the domain of $\psi'(\cdot)$ is compact [14], [20].

C. Gaussian Noise Bound

We consider *i.i.d.* Gaussian noise, denoted by $\mathbf{v}_t \in \mathbb{R}^{n_y}$, for $t = 1, \dots, T$, where $T \in \mathbb{N}$ denotes a finite time horizon. With high probability, jointly throughout all time steps, noise can be bounded by a box, $\subset \mathbb{R}^{n_y}$. Formally:

Lemma 2: Let vectors $\mathbf{v}_1, \dots, \mathbf{v}_T \in \mathbb{R}^{n_y}$, such that for each time $t = 1, \dots, T$ and dimension $j = 1, \dots, n_y$, the noise $v_{t,j} \sim \mathcal{N}(0, \lambda_v^2)$, where $\lambda_v \in \mathbb{R}$. Then, with probability at least $(1 - \delta^v)$, where $\delta^v \in (0, 1)$, jointly for all $t = 1, \dots, T$, we have that $\mathbf{v}_t \in [0 \pm \sqrt{2}\lambda_v \sqrt{\ln(T \cdot n_y / \delta^v)}]^{n_y} \subset \mathbb{R}^{n_y}$.

Proof: The proof is similar to those for Lemma 5.1 in [20] and Lemma 4 in [21]. For each time $t = 1, \dots, T$ and dimension $j = 1, \dots, n_y$, we bound $v_{t,j}$ by applying the Gaussian error function with a probability budget, $\delta^v / (T \cdot n_y)$; we then obtain the result via a union bound over all t and j . \square

III. SYSTEM FORMULATION

We consider a discrete-time dynamical system with finite-horizon $T \in \mathbb{N}$, nonlinear dynamics and observation functions, and additive noises. Formally, for $t = 1, \dots, T$, the system can be described as follows:

$$\begin{aligned} \mathbf{x}_t &= d(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \mathbf{w}_{t-1}) \\ &= f(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) + g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) + \mathbf{w}_{t-1} \\ \mathbf{y}_t &= o(\mathbf{x}_t, \mathbf{u}_t, \mathbf{v}_t) = h(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{v}_t \end{aligned} \quad (1)$$

Here, for all $t = 0, \dots, T$, $\mathbf{x}_t \in \mathbb{R}^{n_x}$, $\mathbf{u}_t \in \mathcal{U} \subset \mathbb{R}^{n_u}$, $\mathbf{y}_t \in \mathbb{R}^{n_y}$, $\mathbf{w}_t \in \mathbb{R}^{n_x}$, and $\mathbf{v}_t \in \mathbb{R}^{n_y}$ denote the system state, control, measurement, process noise, and measurement noise, respectively. We assume that the domain \mathcal{U} is compact, and each element of \mathbf{w}_t and \mathbf{v}_t are *i.i.d.* Gaussian. Formally, for each dimension j , we have $w_{t,j} \sim \mathcal{N}(0, \lambda_w^2)$ and $v_{t,j} \sim \mathcal{N}(0, \lambda_v^2)$, where $\lambda_w, \lambda_v \in \mathbb{R}$. Here, \mathbf{w}_{t+1} and \mathbf{v}_t correspond to \mathbf{v}_t in Section II-B. The function $f(\cdot)$ denotes the prior known dynamic model, which can be a parametric physics model. The functions $g(\cdot)$ and $h(\cdot)$ denote the unknown functions to be learned via GP, and correspond to $\psi(\cdot)$ in Section II-B.

Remark 1: By applying Lemma 2, we construct boxes, denoted by $W \subset \mathbb{R}^{n_x}$ and $V \subset \mathbb{R}^{n_y}$, to bound noises \mathbf{w} and \mathbf{v} , respectively. Formally, given $\delta^w \in (0, 1)$, with probability at least $(1 - \delta^w)$, jointly for all $t = 1, \dots, T$, we have that $\mathbf{w}_{t-1} \in W$. Given $\delta^v \in (0, 1)$, with probability at least $(1 - \delta^v)$, jointly for all $t = 1, \dots, T$, we have that $\mathbf{v}_t \in V$.

IV. PROBLEM DEFINITION

Our goal is to develop a set-based state estimator that can produce *consistent* estimates. Let $\hat{X}_t \subset \mathbb{R}^{n_x}$ denote a set-based state estimate produced by our algorithm at time t . By assuming that the controls are given, the estimation process, at time $t = 1, \dots, T$ can be represented as a recursive function, $\hat{X}_t = E(\hat{X}_{t-1}, \mathbf{u}_{t-1}, \mathbf{u}_t, \mathbf{y}_t)$.

When the dynamic and observation models, $d(\cdot)$ and $o(\cdot)$ (respectively), are known, prior arts in set-based state estimation [12], [13], [22] have achieved *strict* consistency guarantees. In contrast, we focus on a scenario where both $d(\cdot)$ and $o(\cdot)$ are learned with a limited amount of data; hence, we relax the strict consistency and focus on *probabilistic* consistency, or δ -consistency, defined as follows:

Definition 1 (δ -consistency): Given $\delta \in (0, 1)$; an initial set-based estimate, $\hat{X}_0 \subset \mathbb{R}^{n_x}$ such that $\mathbf{x}_0 \in \hat{X}_0$; a sequence of controls, $\{\mathbf{u}_t\}_{t=0}^T \subset \mathcal{U}$; and a sequence of measurements, $\{\mathbf{y}_t\}_{t=1}^T \subset \mathbb{R}^{n_y}$; Then, a state estimator is δ -consistent if the sequence of estimates, $\{\hat{X}_t\}_{t=1}^T$, computed via E , satisfies:

$$\Pr[\forall t = 1 \dots T : \mathbf{x}_t \in \hat{X}_t] \geq 1 - \delta$$

This definition states that δ -consistent estimators are able to guarantee that with high probability, jointly for each time step within a finite time horizon, the set-based estimate contains the true state. Note that: (1) Def. 1 implies that the high-probability consistency guarantee holds *jointly* throughout the finite time horizon, rather than *per time-step* in the form of $\forall t = 1 \dots T : \Pr[\mathbf{x}_t \in \hat{X}_t] \geq 1 - \delta$ [23]. (2) Def. 1 indicates a filtering (rather than smoothing) problem, as future measurements are never used to estimate past or current states.

Problem Statement: *Design a set-based estimator that guarantees δ -consistency under the epistemic uncertainties in the learned models, $g(\cdot)$ and $h(\cdot)$, and the aleatoric uncertainties in the noises, \mathbf{w} and \mathbf{v} .*

V. METHOD

Our recursive set-based estimator, GP-ZKF, shown in Fig. 2, follows the algorithmic structure of Kalman filter and set-based estimators [12], [24]. We choose zonotopes to represent the set-based estimates due to their computational advantages in affine transformation and Minkowski sum (see Section II-A).

At time step $t = 1, \dots, T$, GP-ZKF computes a new zonotopic estimate, $\hat{X}_t \subset \mathbb{R}^{n_x}$ in three phases: **1) Prediction:** Given $\mathbf{u}_{t-1} \in \mathcal{U}$ and $\hat{X}_{t-1} \subset \mathbb{R}^{n_x}$ that contains \mathbf{x}_{t-1} (formally, $\hat{X}_{t-1} \ni \mathbf{x}_{t-1}$), GP-ZKF constructs a *dynamics-consistent* zonotope, $\hat{X}_t \subset \mathbb{R}^{n_x}$, to bound the output of the dynamics, $d(\cdot)$, with high probability. **2) Measurement:** Given $\mathbf{u}_t \in \mathcal{U}$, \hat{X}_t and a new measurement, $\mathbf{y}_t \in \mathbb{R}^{n_y}$, GP-ZKF constructs a *measurement-consistent* polytope, $\bar{X}_{y_t} \subset \mathbb{R}^{n_x}$, that contains all states consistent with \mathbf{y}_t , with high probability. We choose polytopes to represent the potentially asymmetric set, \bar{X}_{y_t} . **3) Correction:** GP-ZKF constructs the new zonotopic estimate, \hat{X}_t , via the intersection $\hat{X}_t \cap \bar{X}_{y_t}$.

A. Phase 1: Prediction

In the prediction phase, given \mathbf{u}_{t-1} and $\hat{X}_{t-1} \ni \mathbf{x}_{t-1}$, GP-ZKF constructs a *dynamics-consistent* zonotope, \hat{X}_t , to bound the outputs of $d(\cdot)$, with high probability. As defined in (1), $d(\cdot)$ is decomposed into three components: a known function, $f(\cdot)$; an unknown but to-be-learned function, $g(\cdot)$; and a noise, \mathbf{w} . In this section, we separately formulate bounds for each component and then integrate them into \hat{X}_{t-1} to bound $d(\cdot)$.

1) Bounding the Known Function: Bounding arbitrary functions is nontrivial. In this work, we focus on smooth functions by making the following assumption:

Assumption 1: i) $f(\cdot)$ is twice continuously differentiable; ii) $f(\cdot)$ is L_f -Lipschitz continuous with respect to the 2-norm; and iii) $\|f(\mathbf{x}_0, \mathbf{u}) - \mathbf{x}_0\|_2 \leq B^f$ for each $\mathbf{x}_0 \in \hat{X}_0$, $\mathbf{u} \in \mathcal{U}$.

In this section, we use Assumption 1(i) to bound the outputs of $f(\cdot)$; we will incorporate Assumption 1(ii, iii) (based on Assumption 3 in [21]) in the following section (Section V-A2).

Assumption 1(i) allows us to directly apply reachability analysis [18] to bound the outputs of $f(\cdot)$, given \mathbf{u}_{t-1} and $\hat{X}_{t-1} \ni \mathbf{x}_{t-1}$. In particular, we first linearize $f(\cdot)$ around a reference point $\bar{\mathbf{x}}_{t-1} \in \mathbb{R}^{n_x}$, which is chosen to be the center of \hat{X}_{t-1} . The linearized function $\bar{f}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) = f(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1}) + J_x^f \cdot (\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1})$, where J_x^f is the Jacobian of $f(\cdot)$ with respect to \mathbf{x}_{t-1} , evaluated at $(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1})$. Given \mathbf{u}_{t-1} and $\hat{X}_{t-1} \ni \mathbf{x}_{t-1}$, the linearization error can be bounded by a zero-centered box that $\subset \mathbb{R}^{n_x}$ (see Prop. 3.7 [18]). Formally, for all $\mathbf{x}_{t-1} \in \hat{X}_{t-1}$,




Fig. 2. A flowchart illustrating the three phases of the set-based estimator, GP-ZKF, at time $t = 1, \dots, T$: (1) Prediction: given the previous zonotopic estimate, \hat{X}_{t-1} , and the control, \mathbf{u}_{t-1} (omitted in the figure), GP-ZKF produces a dynamics-consistent zonotope, \bar{X}_t . (2) Measurement: given \bar{X}_t , the new measurement, \mathbf{y}_t , and the control, \mathbf{u}_t (omitted in the figure), GP-ZKF produces a measurement-consistent polytope, \bar{X}_{y_t} . (3) Correction: GP-ZKF computes the new zonotopic estimate, \hat{X}_t , via $\bar{X}_t \cap \bar{X}_{y_t}$. In the prediction phase, the dynamics contains the known function, $f(\cdot)$; the learned function, $g(\cdot)$; and the process noise, \mathbf{w} (Eq. (1)). In the measurement phase, the observation function contains the learned function, $h(\cdot)$, and the measurement noise, \mathbf{v} (Eq. (2)).

$\mathbf{u}_{t-1} \in \mathcal{U}$, we have:

$$f(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \bar{f}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) \in R^f(\hat{X}_{t-1}, \mathbf{u}_{t-1}) \subset \mathbb{R}^{n_x} \quad (3)$$

Here, $R^f(\cdot)$ denotes the function used to compute the box that bounds the error based on \hat{X}_{t-1} and \mathbf{u}_{t-1} .

2) *Bounding the Unknown Function*: The unknown function $g(\cdot)$ is learned via a GP (see Section II-B). In this section, we formulate a high-probability bound for the output of $g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})$, given \mathbf{u}_{t-1} and $\hat{X}_{t-1} \ni \mathbf{x}_{t-1}$, in the following five steps: i) We state certain regularization assumptions about $g(\cdot)$; ii) show that the state space is compact, with high probability; iii) utilize this compactness to derive a bound for the GP posterior mean, $\mu^g(\cdot)$; iv) present a bound for the GP posterior standard deviation, $\sigma^g(\cdot)$; and (v) integrate both bounds of $\mu^g(\cdot)$ and $\sigma^g(\cdot)$ into the GP confidence intervals, as introduced in Lemma 1, to bound the output of $g(\cdot)$.

(i) *Regularization assumptions*: We denote the kernel for $g(\cdot)$ as k^g and the single-output surrogate function as $g'(\cdot)$ (see Section II-B), and assume the following:

Assumption 2: i) k^g is 2-times continuously differentiable (Def. 4.35 in [25]), ii) k^g is bounded by $\|k^g\|_\infty$ (4.15 in [25]), iii) k^g has bounded derivatives (Assumption 4 in [21]), and (iv) $\|g'\|_{k^g} \leq B^g$ (see Section II-B).

This assumption states that $k^g(\cdot, \cdot)$ and $g(\cdot)$ are smooth and bounded; common smooth kernels, such as square exponential and rational quadratic kernels, satisfy this assumption. Assumption 2 (i, iii, iv) implies that $g(\cdot)$ is L_g -Lipschitz continuous with respect to the 2-norm (Cor. 2 in [21]).

ii) *Compact state space*: The infinite support of the process noise, \mathbf{w} , makes the state space of our system unbounded. Taking a step back, the following lemma states that, with high probability, the *reachable* state space is compact during the estimation processes:

Lemma 3: With Asm. 1 (ii, iii), Asm. 2, and our Gaussian Asm. on \mathbf{w} (Section III), and given initial set $\hat{X}_0 \ni \mathbf{x}_0$, then with probability at least $(1 - \delta^w)$, jointly for all $t = 0, \dots, T$, we have that $\mathbf{x}_t \in \mathcal{X}$, where \mathcal{X} is a box $\subset \mathbb{R}^{n_x}$ (compact) that depends upon $\hat{X}_0, T, \lambda^w, \delta^w, L^f, L^g, B^f, B^g, \|k^g\|_\infty, n_x$.

Lemma 3 can be proven by integrating \hat{X}_0 and our noise bound, W (Remark 1 in Section III), into Lemma 44 in [21]. The intuition behind the proof is that our Assumption 1 (ii, iii), Assumption 2, and the Gaussian nature of \mathbf{w} (with a probability budget δ^w) “prevent” $f(\cdot)$, $g(\cdot)$, and \mathbf{w} (respectively) from deviating the state infinitely far away from \hat{X}_0 .

iii) *Bound posterior mean*: Next, we derive a bound for the posterior mean, $\mu^g(\cdot)$. We first linearize $\mu^g(\cdot)$ around $\bar{\mathbf{x}}_{t-1}$,

which is chosen to be the center of \hat{X}_{t-1} . The linearized function is $\bar{\mu}^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) = \mu^g(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1}) + J_{x_t}^{\mu^g} \cdot (\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1})$, where $J_{x_t}^{\mu^g}$ is the Jacobian of μ^g . Lemma 3 implies that the domain of $\mu^g(\cdot)$ during the estimation process is compact, with high probability. Together with Assumption 2 (i), we obtain that, with high probability, for each $j = 1, \dots, n_x$, the mean μ_j^g is twice continuously differentiable with $L_{\nabla \mu}^g$ -Lipschitz gradient. We then follow the steps in V(A)2) in [14] to derive a bound for the linearization error, as follows:

$$|\mu_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \bar{\mu}_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})| \leq \frac{1}{2} L_{\nabla \mu}^g \cdot \|\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^2 \quad (4)$$

Here, with probability at least $(1 - \delta^w)$, this bound holds uniformly for all $t = 1, \dots, T$, $j = 1, \dots, n_x$, $\mathbf{x}_{t-1} \in \mathcal{X}$, $\mathbf{u}_{t-1} \in \mathcal{U}$, where the probability $(1 - \delta^w)$ is derived from Lemma 3.

iv) *Bound standard deviation*: We approximate the standard deviation, $\sigma^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})$, using $\sigma^g(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1})$. Assumption 2 (i, ii, iii) allows the use of (21,22 in [26]) to bound the approximation error. Formally, with $L_\sigma^g \in \mathbb{R}$, we have for all $t = 1, \dots, T$, $j = 1, \dots, n_x$, $\mathbf{x}_{t-1} \in \mathbb{R}^{n_x}$, $\mathbf{u}_{t-1} \in \mathcal{U}$:

$$|\sigma_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \sigma_j^g(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1})| \leq L_\sigma^g \cdot \|\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^{1/2} \quad (5)$$

v) *Bound combination*: Assumption 2 (iv) allows the use of Lemma 1 to construct confidence intervals for $g(\cdot)$ with $\delta^g \in (0, 1)$. Via a union bound, we combine the confidence intervals and our bounds for the mean and standard deviation to bound the error, $|g(\cdot) - \bar{\mu}^g(\cdot)|$, as follows:

$$\begin{aligned} & |g_j(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \bar{\mu}_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})| \\ & \leq \underbrace{\frac{1}{2} L_{\nabla \mu}^g \cdot \epsilon^2}_{\text{Linearization error of } \mu_j^g(\cdot)} + \underbrace{\beta^g \cdot \sigma_j^g(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1}) + \beta^g \cdot L_\sigma^g \cdot \epsilon^{1/2}}_{\text{Epistemic uncertainty } \oplus \text{ Approx. error of } \sigma_j^g(\cdot)} \quad (6) \end{aligned}$$

Here, $\epsilon = \|\hat{X}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2$ is the norm of the translated zonotope. (We provide the derivation in the Appendix.) Let $R^g(\hat{X}_{t-1}, \mathbf{u}_{t-1})$ be a zero-centered box $\subset \mathbb{R}^{n_x}$, whose radius along each dimension $j = 1, \dots, n_x$ is the RHS of (6). Thus, with probability at least $(1 - \delta^g - \delta^w)$, jointly for all $t = 1, \dots, T$, $\hat{X}_{t-1} \subset \mathcal{X}$, $\mathbf{u}_{t-1} \in \mathcal{U}$, $\mathbf{x}_{t-1} \in \hat{X}_{t-1}$, we have that

$$g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \bar{\mu}^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) \in R^g(\hat{X}_{t-1}, \mathbf{u}_{t-1}) \quad (7)$$

3) *The Prediction Phase*: We define a linear function, $\bar{d}(\cdot)$, by summing the linearized $f(\cdot)$ and $\mu_g(\cdot)$ as $\bar{d}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) = \bar{f}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) + \bar{\mu}^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})$. We define a box, $R^d(\hat{X}_{t-1}, \mathbf{u}_{t-1}) \subset \mathbb{R}^{n_x}$, by Minkowski summing the error

bounds of $f(\cdot)$ (3), $g(\cdot)$ (7), and \mathbf{w} (Remark 1 in Section III). Then, with probability at least $(1 - \delta^g - \delta^w)$, the error between $d(\cdot)$ and $\bar{d}(\cdot)$ can be bounded by $R^d(\cdot)$ as follows:

$$\begin{aligned} d(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \mathbf{w}_{t-1}) - \bar{d}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) &\in R^d(\hat{X}_{t-1}, \mathbf{u}_{t-1}) \\ &:= \underbrace{R^f(\hat{X}_{t-1}, \mathbf{u}_{t-1})}_{\text{Linearization err. of } f(\cdot)} \oplus \underbrace{R^g(\hat{X}_{t-1}, \mathbf{u}_{t-1})}_{\text{Epistemic } \oplus \text{ Lin. err. of } \mu^g(\cdot)} \oplus \underbrace{W}_{\text{Aleatoric}} \oplus \underbrace{\quad}_{\text{Approx. err. of } \sigma^g(\cdot)} \end{aligned} \quad (8)$$

Given that \bar{d} is a linear function, \mathbf{u}_{t-1} is known, and $\mathbf{x}_{t-1} \in \hat{X}_{t-1}$, the output $\bar{d}(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})$ is bounded by a zonotope obtained by linearly transforming \hat{X}_{t-1} . Then, by Minkowski-summing this transformed zonotope and $R^d(\hat{X}_{t-1}, \mathbf{u}_{t-1})$, we obtain another zonotope, denoted by \bar{X}_t , that bounds the output $d(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \mathbf{w}_{t-1})$ (see Section II-A). Let $D(\hat{X}_{t-1}, \mathbf{u}_{t-1})$ denote the function to compute \bar{X}_t ; we now summarize the steps described throughout this section in the following lemma:

Lemma 4: Let $\delta^g, \delta^w \in (0, 1)$. For all $n \in \mathbb{N}$, we choose β_n^g according to Lemma 1. Given an initial set $\hat{X}_{n,0} \ni \mathbf{x}_{n,0}$, then, with probability at least $(1 - \delta^g - \delta^w)$, the following holds jointly for all $n \in \mathbb{N}$, $t = 1, \dots, T$:

- (i) $d(\mathbf{x}_{n,t-1}, \mathbf{u}_{n,t-1}, \mathbf{w}_{n,t-1}) \in \bar{X}_{n,t} = D(\hat{X}_{n,t-1}, \mathbf{u}_{n,t-1})$, for all $\hat{X}_{n,t-1} \subset \mathcal{X}$, $\mathbf{x}_{n,t-1} \in \hat{X}_{n,t-1}$, $\mathbf{u}_{n,t-1} \in \mathcal{U}$, where $\mathbf{w}_{n,t-1}$ is the process noise as assumed in Section III.
- (ii) $\mathbf{x}_{n,t} \in \mathcal{X}$, $\mathbf{x}_{n,0} \in \mathcal{X}$.

Proof: The noise bound, W (Remark 1), and Lemma 3 hold jointly, with probability at least $(1 - \delta^w)$. Then, by applying a union bound to combine the above result with the confidence intervals of $g(\cdot)$ (Lemma 1), we arrive at the bound, $\bar{X}_{n,t}$. \square

Hence, the prediction phase computes the dynamic-consistent zonotope, $\bar{X}_t = D(\hat{X}_{t-1}, \mathbf{u}_{t-1}) \subset \mathbb{R}^{n_x}$, which bounds the output of the dynamic function, $d(\cdot)$ (Lemma 4).

B. Phase 2: Measurement

In the measurement phase, given \bar{X}_t from Lemma 4 and \mathbf{u}_t , GP-ZKF constructs a *measurement-consistent* polytope [24], \bar{X}_{y_t} , to bound all possible states consistent with the new measurement, \mathbf{y}_t , with high probability.

The unknown function $h(\cdot)$ is learned via GP (see Section II-B). We denote the single-output surrogate function by $h'(\cdot)$, and the kernel by k^h . Similar to Assumption 2 about k^g and $g(\cdot)$, we make the following assumption about k^h and $h(\cdot)$:

Assumption 3: (i) k^h is 2-times continuously differentiable (Def. 4.35 in [25]), (ii) k^h is bounded by $\|k^h\|_\infty$ (4.15 in [25]), (iii) k^h has bounded derivatives (Assumption 4 in [21]), and (iv) $\|h'\|_{k^h} \leq B^h$ (see Section II-B).

To bound $h(\cdot)$, we follow the same steps described in Section V-A2 to bound $g(\cdot)$. **Bound posterior mean:** We linearize $\mu^h(\cdot)$ around $\bar{\mathbf{x}}_t$ = the center of \bar{X}_t , and obtain $\bar{\mu}^h(\mathbf{x}_t, \mathbf{u}_t) = \mu^h(\bar{\mathbf{x}}_t, \mathbf{u}_t) + J_x^{\mu^h} \cdot (\mathbf{x}_t - \bar{\mathbf{x}}_t)$. Lemma 4 (ii) implies the compactness of the domain of $\mu^h(\cdot)$ within the estimation process, which leads to a linearization error bound (similar to (4)). **Bound standard deviation:** With Assumption 3, we derive a bound for $\sigma^h(\cdot)$ (similar to (5)). **Bound combination:** Given $\delta^h \in (0, 1)$, Lemma 1 allows us to construct confidence intervals for $h(\cdot)$. Via a union bound, we determine that the confidence intervals and noise bound $\mathbf{v}_t \in V$ (see Remark 1) jointly hold with probability at least $(1 - \delta^h - \delta^v)$. Then, similar to R^g in (7),

we obtain a box, $R^h(\bar{X}_t, \mathbf{u}_t) \subset \mathbb{R}^{n_y}$. Thus, with probability at least $(1 - \delta^g - \delta^w)(1 - \delta^h - \delta^v)$, jointly for all $t = 1, \dots, T$, $\bar{X}_t \subset \mathcal{X}$, $\mathbf{u}_t \in \mathcal{U}$, $\mathbf{x}_t \in \bar{X}_t$, we have that

$$h(\mathbf{x}_t, \mathbf{u}_t) - \bar{\mu}^h(\mathbf{x}_t, \mathbf{u}_t) \in R^h(\bar{X}_t, \mathbf{u}_t) \quad (9)$$

Here, the product rule, $(1 - \delta^g - \delta^w)(1 - \delta^h - \delta^v)$, results from the assumption that noises \mathbf{w} and \mathbf{v} are independent with each other (Section III). By expanding $\bar{\mu}^h$ and then combining the noise bound V (Remark 1) with (9), we obtain that with probability at least $(1 - \delta^g - \delta^w)(1 - \delta^h - \delta^v)$, jointly for all $t = 1, \dots, T$, $\bar{X}_t \subset \mathcal{X}$, $\mathbf{u}_t \in \mathcal{U}$, $\mathbf{x}_t \in \bar{X}_t$, the following holds:

$$\begin{aligned} &\bar{\mu}^h(\mathbf{x}_t, \mathbf{u}_t) - o(\mathbf{x}_t, \mathbf{u}_t, \mathbf{v}_t) \\ &= \mu^h(\bar{\mathbf{x}}_t, \mathbf{u}_t) + J_x^{\mu^h} \cdot (\mathbf{x}_t - \bar{\mathbf{x}}_t) - h(\mathbf{x}_t, \mathbf{u}_t) - \mathbf{v}_t \\ &\in R^o(\bar{X}_t, \mathbf{u}_t) := \underbrace{R^h(\bar{X}_t, \mathbf{u}_t)}_{\text{Epistemic } \oplus \text{ Lin. err. of } \mu^h(\cdot)} \oplus \underbrace{V}_{\text{Aleatoric}} \oplus \underbrace{\quad}_{\text{Approx. err. of } \sigma^h(\cdot)} \end{aligned} \quad (10)$$

Given a new measurement, $\mathbf{y}_t = o(\mathbf{x}_t, \mathbf{u}_t, \mathbf{v}_t) \in \mathbb{R}^{n_y}$ (2), the measurement-consistent states must satisfy the bound in (10). Hence, we equivalently represent the bound in (10) as a polytope, $\bar{X}_{y_t} \subset \mathbb{R}^{n_x}$, with \mathbf{x}_t as the variable, as follows:

$$\begin{aligned} \bar{X}_{y_t} = \{ \mathbf{x}_t \in \mathbb{R}^{n_x} : \\ J_x^{\mu^h} \cdot \mathbf{x}_t - [\mathbf{y}_t - \mu^h(\bar{\mathbf{x}}_t, \mathbf{u}_t) + J_x^{\mu^h} \cdot \bar{\mathbf{x}}_t] \in R^o(\bar{X}_t, \mathbf{u}_t) \} \end{aligned} \quad (11)$$

The output of the measurement step is the measurement-consistent polytope, $\bar{X}_{y_t} \subset \mathbb{R}^{n_x}$. Let $O^{inv}(\bar{X}_t, \mathbf{u}_t, \mathbf{y}_t)$ denote the function to compute \bar{X}_{y_t} , where the superscript *inv* emphasizes that O^{inv} is the ‘‘inverse’’ of our observation model, $o(\cdot)$.

C. Phase 3: Correction

In the correction phase, the goal is to construct a zonotope \hat{X}_t as the intersection $\bar{X}_t \cap \bar{X}_{y_t}$, where \bar{X}_t is formulated in Lemma 4 and \bar{X}_{y_t} is defined in (11). Note that the bound in (10) requires $\bar{X}_t \subset \mathcal{X}$; therefore, GP-ZKF instead computes the intersection $\bar{X}_t \cap \bar{X}_{y_t} \cap \mathcal{X}$. This intersection cannot be computed exactly, and hence is outer-approximated by \hat{X}_t . The new zonotope, \hat{X}_t , is potentially less conservative than \bar{X}_t and \bar{X}_{y_t} in bounding the possible states, as \hat{X}_t takes into account both dynamics and measurements. Formally:

Lemma 5: Let $\delta^g, \delta^w, \delta^h, \delta^v \in (0, 1)$. For all $n^g, n^h \in \mathbb{N}$, we choose $\beta_{n^g}^g, \beta_{n^h}^h$ according to Lemma 1¹. Given an initial set $\hat{X}_0 \ni \mathbf{x}_0$, then, with probability at least $(1 - \delta^g - \delta^w)(1 - \delta^h - \delta^v)$, jointly for all $n^g, n^h \in \mathbb{N}$, $t = 1, \dots, T$, $\hat{X}_{t-1} \subset \mathcal{X}$, $\mathbf{x}_{t-1} \in \hat{X}_{t-1}$, $\mathbf{u}_{t-1}, \mathbf{u}_t \in \mathcal{U}$, and $\mathbf{y}_t \in \mathbb{R}^{n_y}$, we have the following:

$$d(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \mathbf{w}_{t-1}) \in (\bar{X}_t \cap \bar{X}_{y_t} \cap \mathcal{X}) \subset \hat{X}_t \quad (12)$$

where $\mathbf{w}_{n,t-1}$ is the process noise as assumed in Section III, $\bar{X}_t = D(\hat{X}_{t-1}, \mathbf{u}_{t-1})$ (Lemma 4), and $\bar{X}_{y_t} = O^{inv}(\bar{X}_t, \mathbf{u}_t, \mathbf{y}_t)$ (11).

This lemma states that the true state at time t lies within the zonotope, \hat{X}_t , with high probability. Lemma 5 summarizes the derivations in Section V-B; it can be proved by directly combining Lemma 4 and the bound in (10).

¹We omit the subscripts n^g, n^h for every variable in this lemma for clarity.

TABLE I
 THE ESTIMATION RESULTS IN THE SIMULATED PENDULUM DOMAIN

Data shift	Method	RMSE ($\theta, \dot{\theta}$) ($^\circ, ^\circ/s$)	Incl. (%)	Radius ($\theta, \dot{\theta}$) ($^\circ, ^\circ/s$)	Time (s)
Both models	GP-EKF	20.2, 37.5	5	1.7, 9.2	0.003
	GP-UKF	7.4, 15.6	38	10.0, 32.1	0.009
	GP-PF	90.5, 64.4	28	27.0, 54.6	1.459
	GP-ZKF	16.4, 20.4	83	46.7, 131.4	0.004
Dynamics model	GP-EKF	15.3, 22.1	0.17	0.2, 0.4	0.003
	GP-UKF	2.2, 10.8	27	1.6, 17.8	0.009
	GP-PF	69.6, 46.9	16	11.8, 25.5	1.500
	GP-ZKF	16.4, 21.5	88	44.6, 72.5	0.004
Observation model	GP-EKF	15.6, 21.7	4	1.4, 2.0	0.003
	GP-UKF	25.4, 31.2	27	35.9, 40.6	0.009
	GP-PF	222.7, 157.4	18	82.9, 100.2	1.544
	GP-ZKF	15.3, 18.6	87	47.9, 126.0	0.004
None	GP-EKF	15.4, 20.4	0.00	0.2, 0.4	0.003
	GP-UKF	2.2, 7.8	18	1.5, 5.1	0.009
	GP-PF	161.3, 126.7	10	47.0, 47.8	1.573
	GP-ZKF	16.6, 19.7	92	47.1, 72.8	0.004

The intersection in (12) is outer-approximated in two steps: First, GP-ZKF follows Prop. 1 in [24] to obtain a zonotope denoted by $Z_t(\Lambda_t)$ and parameterized by the matrix Λ_t , such that $Z_t(\Lambda_t) \supset (\hat{X}_t \cap \bar{X}_{y_t})$. The parameter Λ_t is obtained by analytically solving a convex program that minimizes the “size” of $Z_t(\Lambda_t)$ (see Sec. 6.1 in [12]). Then, GP-ZKF follows the same procedures to construct $\hat{X}_t \supset (\bar{X}_t \cap \bar{X}_{y_t} \cap \mathcal{X})$.

The essence of the correction step can be represented by the function $E(\hat{X}_{t-1}, \mathbf{u}_{t-1}, \mathbf{u}_t, \mathbf{y}_t)$, which conducts the intersections and outputs the new zonotopic estimate, \hat{X}_t .

VI. THEORETICAL GUARANTEES

We present two key theorems about GP-ZKF. First, as both epistemic and aleatoric uncertainties have been bounded, GP-ZKF is δ -consistent (Def. 1); second, when both uncertainties are relaxed, GP-ZKF reduces to GP-EKF [7].

Theorem 1: Given $\delta \in (0, 1)$, GP-ZKF chooses $\delta^g, \delta^h, \delta^w, \delta^v \in (0, 1)$ such that $(1 - \delta^g - \delta^w)(1 - \delta^h - \delta^v) \geq (1 - \delta)$. For all $n^g, n^h \in \mathbb{N}$, GP-ZKF chooses $\beta_{n^g}^g, \beta_{n^h}^h$ according to Lemma 1. Then, GP-ZKF is δ -consistent.

Proof: Similar to the proof for Cor. 7 in [14], we recursively apply Lemma 5 from $t = 1$ to T and obtain with high probability, jointly for all $t = 1, \dots, T$, that $\mathbf{x}_t \in \hat{X}_t$. \square

GP-EKF [7] is an EKF-based state estimator that learns both dynamic and observation models via GPs; we see GP-EKF as the *stochastic* counterpart to our *set-based* GP-ZKF. At every time $t = 1, \dots, T$, GP-EKF computes the Kalman gain, $K_t \in \mathbb{R}^{n_x \times n_y}$, and outputs a point estimate, $\mu_t \in \mathbb{R}^{n_x}$, and a covariance, $\Sigma_t \in \mathbb{R}^{n_x \times n_x}$ (see Table II in [7]). To draw an analogy similar to Thm.7 in [16], we define GP-ZKF’s Kalman gain as the matrix Λ_t , which parameterizes the intersection $\bar{X}_t \cap \bar{X}_{y_t}$, as mentioned at the end of Section V-C. We define GP-ZKF’s point estimate as the zonotope center, $(\hat{X}_t)_c$, and covariance as $(\hat{X}_t)_G((\hat{X}_t)_G)^T$ (the zonotope covariation, as seen in Def. 4 in [16]). Next, we formally demonstrate that GP-ZKF can be reduced to (the Joseph form of) GP-EKF:

Theorem 2: Assume that GP-ZKF’s set-based estimates are always inside the state space—i.e., for all $t = 1, \dots, T$, the zonotope $Z_t(\Lambda_t) \subset \mathcal{X}$, where $Z_t(\Lambda_t)$ is defined at the end of




Fig. 3. Zonotopic estimates along a trajectory produced by GP-ZKF under the *Shift None* condition. The zonotopic estimates, \hat{X}_t (green fill), always outer-approximate the intersection between the dynamic-consistent zonotope, \bar{X}_t (yellow fill), and the measurement-consistent zonotope, \bar{X}_{y_t} (blue outline). (1) Despite the erroneous point estimates, defined as $(\hat{X}_t)_c$ (green dots), the zonotopes, \hat{X}_t , still contain the true states (black dots). (2) Even though the volumes of \bar{X}_t may grow rapidly during uncertainty propagation, the volumes of \hat{X}_t shrink when informative measurements are available.

Section V-C. Given the same initial condition, $\mu_0 = (\hat{X}_0)_c$ and $K_0 = (\hat{X}_0)_G((\hat{X}_0)_G)^T$, if GP-ZKF: (i) sets the confidence interval scalings for $g(\cdot)$ and $h(\cdot)$, denoted by β^g and β^h (respectively), to 1; (ii) omits all noise bounds for \mathbf{w} and \mathbf{v} by setting $W = \emptyset$ and $V = \emptyset$ (Remark 1); and (iii) omits all linearization errors for $f(\cdot)$, μ^g , σ^g , μ^h , and σ^h by setting $R^f(\cdot) = \emptyset$ (3), $L_{\nabla\mu}^g = 0$ (4), $L_{\sigma}^g = 0$ (5), $L_{\nabla\mu}^h = 0$ (Section V-B), and $L_{\sigma}^h = 0$ (Section V-B); then, for time $t = 1, \dots, T$, we have $K_t = \Lambda_t$, $\mu_t = (\hat{X}_t)_c$, and $\Sigma_t = (\hat{X}_t)_G((\hat{X}_t)_G)^T$.

Proof: Under the relaxations above, (12) becomes $(\bar{X}_t \cap \bar{X}_{y_t} \cap \mathcal{X}) \subset Z_t(\Lambda_t) = \hat{X}_t$. And each of the uncertainty bounds, $R^d(\cdot)$ (8) and $R^o(\cdot)$ (10), only contains one standard deviation. As introduced at the end of Section V-C, GP-ZKF obtains $\hat{X}_t(\Lambda_t)$ by optimizing Λ_t . With the analytical solution, Λ_t , we reach the final conclusion by induction (see the proof of Thm.7 in [16]). \square

This theorem states that with certain relaxations, GP-ZKF could produce the same Kalman gain, point estimate, and covariance as GP-EKF. The Kalman gain in GP-ZKF, Λ_t , weighs the dynamic-consistent zonotope, \bar{X}_t , and the measurement-consistent polytope, \bar{X}_{y_t} , when “mixing” them within the outer-approximated intersection. In contrast to GP-EKF, Thm. 2 signifies the conservativeness of GP-ZKF in bounding the linearization errors and aleatoric and epistemic uncertainties during estimation. The conservativeness echoes GP-ZKF’s consistency guarantee, as stated in Thm. 1.

VII. EXPERIMENT AND RESULT

In Section VI, we theoretically highlight GP-ZKF’s consistency guarantee and its connection with GP-EKF. We also empirically demonstrate GP-ZKF’s improved consistency against three GP-based stochastic approaches: GP-EKF, GP-UKF, and GP-PF [7], in both a simulated inverted pendulum domain under significant epistemic uncertainty and a real-world robot-assisted dressing domain.

TABLE II
THE ESTIMATION RESULTS IN THE ROBOT-ASSISTED DRESSING DOMAIN

Arm pose	Method	RMSE (x, y, z) (cm,cm,cm)	Incl. (%)	Radius (x, y, z) (cm,cm,cm)	Time (s)
Bend	GP-EKF	3.9, 3.3, 3.8	83	11, 19, 23	0.25
	GP-UKF	9.8, 3.0, 6.1	76	26, 10, 15	1.68
	GP-PF	2.4, 2.6, 3.0	73	6, 6, 7	29.93
	GP-ZKF	2.8, 3.4, 4.3	88	8, 8, 8	0.25
Lower	GP-EKF	5.2, 3.0, 6.9	94	24, 19, 22	0.11
	GP-UKF	9.2, 12.4, 13.6	80	27, 22, 33	0.71
	GP-PF	5.0, 2.7, 4.1	70	11, 8, 11	12.74
	GP-ZKF	4.6, 2.7, 4.7	97	12, 10, 12	0.11
Straight	GP-EKF	2.0, 3.0, 3.9	88	14, 15, 14	0.12
	GP-UKF	3.4, 3.8, 3.6	64	7, 7, 9	0.78
	GP-PF	2.5, 3.0, 2.8	62	6, 6, 6	13.95
	GP-ZKF	1.7, 2.8, 4.1	92	8, 8, 8	0.12

In both domains, we evaluated all methods using the following metrics: 1) Avg. root-mean-square error (RMSE) (per dimension) for the point estimate. GP-ZKF’s point estimate is defined as the zonotope center, $(\hat{X}_t)_c$ (see Section VI). 2) Inclusion: the percentage of time steps during which the true state is “included” by the set-based estimate, which measures a method’s consistency. Recall that Thm. 2 draws equivalence between GP-ZKF’s zonotopic estimate and the *unscaled* covariance from GP-EKF. In practice, we defined GP-EKF, GP-UKF, and GP-PF’s set-based estimates as the 95% confidence ellipsoids determined by the *up-scaled* covariance matrices (GP-PF’s posteriors were approximated as Gaussians) such that they can be on par with GP-ZKF. 3) Avg. radius (per dimension): the radius of the box that outer-approximates the set-based estimates, which measures conservativeness. 4) Avg. computation time per time step.

A. Simulated Pendulum Domain

We simulated a discrete-time 2D inverted pendulum controlled by an infinite-horizon linear quadratic regulator. The state, $\mathbf{x} = [\theta, \dot{\theta}]^T$ (angle and angular velocity), and the set-point correspond to the pendulum standing upright (0°). The dynamics, $d(\cdot)$, is the closed-loop system, with noise \mathbf{w} ($\lambda_w = 7.16^\circ$). The observation function, $o(\cdot)$, maps the state to the end effector position and velocity $\in \mathbb{R}^4$, with noise \mathbf{v} ($\lambda_v = 8.88^\circ$). The known dynamics, $f(\cdot)$, is given by the linearized and discretized system around the set-point.

To evaluate the methods’ consistency under significant epistemic uncertainties, we induced distribution shifts between training and testing. For testing, we ran each method with $T = 15$ from four uniformly sampled starting points within the *testing region*, $\theta \in [\pi, 2\pi]$, with 10 repetitions per starting point. The training data included four variants: 1) *Shift Both*: the GPs for both $g(\cdot)$ and $h(\cdot)$ were trained with a *default* dataset with nine rollouts inside the *training region*, $\theta \in [0, \pi]$. The epistemic uncertainties in both the dynamic and observation models were high due to the shift from the training region to the testing region. 2) *Shift Dynamics*: the GP for $h(\cdot)$ was trained with an additional dataset whose input locations were uniformly sampled in the *testing region*. 3) *Shift Observation*: the GP for $g(\cdot)$ was trained with five additional rollouts inside the *testing region*. 4) *Shift None*: $g(\cdot)$ and $h(\cdot)$ were each trained with their own additional dataset besides the default dataset. Similar to [14], confidence interval scalings β^g and β^h are conservative. We specified and




Fig. 4. Keyframes of the robot-assisted dressing task (cloth is omitted). The simulated robot and human motions were reproduced based on real-world data (Fig. 1). Each keyframe illustrates a few ground-truth states of the human arm and the zonotopic estimate of the human elbow position (outer-approximated as boxes $\subset \mathbb{R}^3$, visualized in green).

scaled them when more data was available, to reflect their dependency upon the information capacity, γ .

We present our results in Table I. Due to the nonlinearity of the dynamic and observation functions, the linearization-based method (GP-EKF) performed poorly, with very low inclusions and small radii. GP-UKF (not linearization-based) performed better, with lower RMSEs, higher inclusions and larger radii. GP-PF produced large RMSEs, potentially due to the narrowness of GP’s predictive distributions. In contrast, GP-ZKF achieved the highest amount of inclusions for all conditions; this empirically validates GP-ZKF’s conservativeness in respecting linearization errors and epistemic and aleatoric uncertainties, which echos GP-ZKF’s probabilistic consistency guarantee (Thm. 1). Fig. 3 illustrates the zonotopic estimates produced by GP-ZKF under the *Shift None* condition. Admittedly, GP-ZKF is more conservative than the others, resulting in larger radii. We argue that GP-ZKF’s conservativeness actually scales appropriately with the domain, as we will next demonstrate its low conservativeness in the dressing domain.

B. Robot-Assisted Dressing Domain

A robot arm was controlled to dress a *long-sleeved* jacket onto a human arm (Fig. 1). We evaluated all methods offline to estimate the visually occluded positions of human elbow [1].

During data collection, the human naturally moved his arm, the configuration of which was tracked by the Xsens motion capture system, which does not suffer from visual occlusion. The robot was controlled to move from the human hand to the elbow, and then to the shoulder position. The data is composed of three initial arm conditions: *bend*, *lower*, and *straight*, with 17, 11, and 12 trajectories, respectively.

The state, $\mathbf{x} \in \mathbb{R}^3$, is the human elbow position. The control, $\mathbf{u} \in \mathbb{R}^9$, contains the positions of the human hand and shoulder, and the robot end effector. The measurement $\mathbf{y} \in \mathbb{R}^3$ is the processed force signal, obtained at the robot end effector. We applied a low-pass filter to the force and converted it to the approximate position of the center of the cuff via a tether-inspired parametric model [27]. The known dynamic model is specified as $f(\mathbf{x}, \mathbf{u}) = \mathbf{x}$; the variances of the noises, \mathbf{w} and \mathbf{v} , are automatically identified within GP regression.

We present the cross-validated results within each initial arm condition in Table II. Both GP-EKF and GP-UKF achieved high amount of inclusions that were comparable with GP-ZKF; however, GP-ZKF’s radii were much smaller, demonstrating its appropriate conservativeness. Essentially, GP-ZKF provided consistent set-based estimates without being overly conservative; Fig. 4 illustrates the zonotopic estimates it produced.

VIII. CONCLUSION AND FUTURE WORK

This work proposes a set-based estimator, GP-ZKF, that is able to produce zonotopic estimates based on learned dynamic and observation models. Theoretically, GP-ZKF guarantees probabilistic consistency. In addition, the *stochastic* approach, GP-EKF, can be seen as a special case of our *set-based* method, GP-ZKF—*i.e.*, where linearization errors and aleatoric uncertainties are omitted and epistemic uncertainties are simplified. Empirically, GP-ZKF outperformed the *stochastic* baselines (GP-EKF, GP-UKF, and GP-PF [7]) in both a simulated pendulum and real-world dressing domain. Future work will focus on combining our approach with control methods [28], [29], as well as richer and more computationally scalable models for force-based estimation.

APPENDIX

We combine the confidence intervals for $g(\cdot)$ (Lemma 1) with the bounds for μ^g (4) and σ^g (5) to bound the error $|g_j(\cdot) - \bar{\mu}_j^g(\cdot)|$ for each $j = 1, \dots, n_x$ as follows:

$$\begin{aligned} & |g_j(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \bar{\mu}_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})| \\ & \leq |\mu_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \bar{\mu}_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})| \\ & \quad + |g_j(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) - \mu_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})| \quad (b) \\ & \leq \frac{1}{2} L_{\nabla \mu}^g \cdot \|\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^2 + \beta^g \cdot \sigma_j^g(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) \quad (c) \\ & \leq \frac{1}{2} L_{\nabla \mu}^g \cdot \|\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^2 \\ & \quad + \beta^g \cdot \sigma_j^g(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1}) + \beta^g L_{\sigma}^g \cdot \|\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^{1/2} \quad (d) \\ & \leq \frac{1}{2} L_{\nabla \mu}^g \cdot \|\hat{\mathbf{X}}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^2 \\ & \quad + \beta^g \cdot \sigma_j^g(\bar{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1}) + \beta^g L_{\sigma}^g \cdot \|\hat{\mathbf{X}}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2^{1/2} \quad (e) \end{aligned}$$

Here, (b) is obtained via the triangle inequality, and we achieve (c) by using a union bound to combine the mean's error bound (4) and confidence intervals (Lemma 1). Hence, all inequalities starting at (c) hold with probability at least $(1 - \delta^g - \delta^w)$. (d) is obtained by applying the standard deviation's error bound (5); (e) is due to the norm of the translated zonotope $\|\hat{\mathbf{X}}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2 := \max_{\mathbf{x}_{t-1} \in \hat{\mathbf{X}}_{t-1}} \|\mathbf{x}_{t-1} - \bar{\mathbf{x}}_{t-1}\|_2$ [30].

REFERENCES

- [1] F. Zhang, A. Cully, and Y. Demiris, "Probabilistic real-time user posture tracking for personalized robot-assisted dressing," *IEEE Trans. Robot.*, vol. 35, no. 4, pp. 873–888, Aug. 2019.
- [2] A. Xu and G. Dudek, "Optimo: Online probabilistic trust inference model for asymmetric human-robot collaborations," in *Proc. 10th ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2015, pp. 221–228.
- [3] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Trust-aware decision making for human-robot collaboration: Model learning and planning," *ACM Trans. Hum.-Robot Interact.*, vol. 9, no. 2, pp. 1–23, 2020.
- [4] R. Choudhury, G. Swamy, D. Hadfield-Menell, and A. D. Dragan, "On the utility of model learning in HRI," in *Proc. 14th ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2019, pp. 317–325.
- [5] V. V. Unhelkar, S. Li, and J. A. Shah, "Decision-making for bidirectional communication in sequential human-robot collaborative tasks," in *Proc. 15th ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2020, pp. 329–341.

- [6] J. A. Castellanos, J. Neira, and J. D. Tardós, "Limits to the consistency of EKF-based SLAM," *IFAC Proc. Volumes*, vol. 37, no. 8, pp. 716–721, 2004.
- [7] J. Ko and D. Fox, "GP-bayesfilters: Bayesian filtering using gaussian process prediction and observation models," *Auton. Robots*, vol. 27, no. 1, pp. 75–90, 2009.
- [8] S. J. Julier and J. K. Uhlmann, "A counter example to the theory of simultaneous localization and map building," in *Proc. Int. Conf. Robot. Automat.*, 2001, pp. 4238–4243.
- [9] G. P. Huang, A. I. Mourikis, and S. I. Roumeliotis, "Analysis and improvement of the consistency of extended Kalman filter based SLAM," in *Int. Conf. Robot. Automat.*, 2008, pp. 473–479.
- [10] S. Huang and G. Dissanayake, "Convergence and consistency analysis for extended Kalman filter based SLAM," *IEEE Trans. Robot.*, vol. 23, no. 5, pp. 1036–1049, Oct. 2007.
- [11] A. Barrau and S. Bonnabel, "An EKF-SLAM algorithm with consistency properties," 2015, *arXiv:1510.06263*.
- [12] T. Alamo, J. M. Bravo, and E. F. Camacho, "Guaranteed state estimation by zonotopes," *Automatica*, vol. 41, no. 6, pp. 1035–1043, 2005.
- [13] B. S. Rego, G. V. Raffo, J. K. Scott, and D. M. Raimondo, "Guaranteed methods based on constrained zonotopes for set-valued state estimation of nonlinear discrete-time systems," *Automatica*, vol. 111, 2020, Art. no. 108614.
- [14] T. Koller, F. Berkenkamp, M. Turchetta, J. Boedecker, and A. Krause, "Learning-based model predictive control for safe exploration and reinforcement learning," in *Proc. IEEE Conf. Decis. Control*, 2019, pp. 6059–6066.
- [15] C. Combastel, "An extended zonotopic and Gaussian Kalman filter (EZGKF) merging set-membership and stochastic paradigms: Toward non-linear filtering and fault detection," *Annu. Rev. Control*, vol. 42, pp. 232–243, 2016.
- [16] C. Combastel, "Zonotopes and Kalman observers: Gain optimality under distinct uncertainty paradigms and robust convergence," *Automatica*, vol. 55, pp. 265–273, 2015.
- [17] J. K. Scott, D. M. Raimondo, G. R. Marseglia, and R. D. Braatz, "Constrained zonotopes: A new tool for set-based estimation and fault detection," *Automatica*, vol. 69, pp. 126–136, 2016.
- [18] M. Althoff, "Reachability analysis and its application to the safety assessment of autonomous cars," Ph.D. dissertation, Technische Universität München, München, Germany, 2010.
- [19] I. Bogunovic, A. Krause, and J. Scarlett, "Corruption-tolerant gaussian process bandit optimization," in *Int. Conf. Artif. Intell. Statist.*, 2020, pp. 1071–1081.
- [20] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for Gaussian process optimization in the bandit setting," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3250–3265, May 2012.
- [21] F. Berkenkamp, "Safe exploration in reinforcement learning: Theory and applications in robotics," Ph.D. dissertation, ETH Zurich, Zürich, Switzerland, 2019.
- [22] C. Combastel, "A state bounding observer for uncertain non-linear continuous-time systems based on zonotopes," in *Proc. 44th IEEE Conf. Decis. Control*, 2005, pp. 7228–7234.
- [23] A. Shetty and G. X. Gao, "Predicting state uncertainty bounds using non-linear stochastic reachability analysis for urban GNSS-based UAS navigation," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 9, pp. 5952–5961, Sep. 2021.
- [24] V. T. H. Le, C. Stoica, T. Alamo, E. F. Camacho, and D. Dumur, "Zonotope-based set-membership estimation for multi-output uncertain systems," in *Proc. IEEE Int. Symp. Intell. Control*, 2013, pp. 212–217.
- [25] I. Steinwart and A. Christmann, *Support Vector Machines*. Berlin, Germany: Springer, 2008.
- [26] A. Lederer, J. Umlauf, and S. Hirche, "Uniform error bounds for Gaussian process regression with application to safe control," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 659–669.
- [27] M. Tognon, R. Alami, and B. Siciliano, "Physical human-robot interaction with a tethered aerial vehicle: Application to a force-based human guiding problem," *IEEE Trans. Robot.*, vol. 37, no. 3, pp. 723–734, Jun. 2021.
- [28] S. Li, N. Figueroa, A. Shah, and J. A. Shah, "Provably safe and efficient motion planning with uncertain human dynamics," in *Proc. Robot.: Sci. Syst.*, 2021.
- [29] T. Stouraitis, I. Chatzinikolaidis, M. Gienger, and S. Vijayakumar, "Online hybrid motion planning for dyadic collaborative manipulation via bilevel optimization," *IEEE Trans. Robot.*, vol. 36, no. 5, pp. 1452–1471, Oct. 2020.
- [30] M. Althoff, "Cora 2021 manual," 2021.