Probabilistically Valid Stochastic Extensions of Deterministic Models for Systems with Uncertainty

Konstantinos Karydis, Ioannis Poulakakis, Jianxin Sun, and Herbert G. Tanner

***Abstract*—Models capable of capturing and reproducing the variability observed in experimental trials can be valuable for planning and control in the presence of uncertainty. This paper reports on a new data-driven methodology that extends deterministic models to a stochastic regime and offers probabilistic guarantees. From an acceptable deterministic model, a stochastic one is generated, capable of capturing and reproducing uncertain system-environment interactions at given levels of fidelity. The reported approach combines methodological elements of probabilistic model validation and randomized algorithms, to simultaneously quantify the fidelity of a model and tune the distribution of random parameters in the augmented stochastic extension, in order to reproduce the variability observed experimentally in a physical process of interest. The approach can apply to an array of physical processes, the models of which may come in different forms, including differential equations; we demonstrate this point by considering examples from the areas of miniature legged robots and aerial vehicles.**

***Index terms*—** Stochastic Extensions, Probabilistic Validation, Model Fidelity, Uncertain Systems, Miniature Legged Robots, Aerial Robots

I. Introduction

Robot control algorithms are predominantly model-based, and often a large part of the effort prior to deployment is devoted to deriving and validating models that can faithfully represent robot behavior. By their very nature, robots interact physically with their environment, and in field deployment these interactions become increasingly uncertain. Examples include vehicles operating in partially-known, dynamic environments (Aoude et al, 2013); legged robots moving on rough terrain (Shkolnik et al, 2011) or fluidizing ground (Qian et al, 2012); quadrotors flying under the influence of uncertain aerodynamic effects (Powers et al, 2012; Zarovy et al, 2010); underwater robots affected by uncertain ocean currents (Pereira et al, 2013); and steerable needles interacting with soft tissue (Alterovitz et al, 2008).

In many of these examples, deterministic models have limited ability to predict the behavior of the robot as it operates in its environment (Chirikjian, 2012; Thrun et al, 2005; Timcenko and Allen, 1993). Still, there is benefit in quantifying the degree of uncertainty in these models. In this paper, we describe a new framework that takes an underlying deterministic model and extends it into a stochastic regime. The augmented stochastic model can then be used to characterize—in a probabilistic sense—the capability of the model to capture the behavior of a system in the presence of uncertainty. Our approach relies on experimental measurements from the physical system, and quantifies the extent to which the stochastic extension captures and reproduces—within a user-mandated confidence level—the range of the behaviors observed in experiments.

The main idea is to parameterize appropriately an otherwise adequate deterministic model of the system, to produce an augmented *stochastic* model. Then, randomized algorithms (Tempo et al, 2012; Vidyasagar, 2001, 2003) can be used to quantify the extent to which the resulting stochastic model captures the uncertain system-environment interactions. In particular, our method hinges on the concept of checking parameterized *distributions of models* against available experimental data. The probabilistic validation part involves a Monte Carlo simulation for estimating the probability that a random model instantiation is statistically consistent with the measurements. Randomized optimization (Vidyasagar, 2001) can then provide approximate near optima for valid model parameters. In this way, data variability is integrated within, and can be reproduced by the model. Essentially, data statistics are used to quantify the amount of the uncertainty that the model parameters need to have to capture the variability observed in the experimental data.

This methodology is general enough to accommodate different robot platforms and types of models. For the shake of concreteness, we demonstrate here its application on two classes of systems that are representative examples of robotic systems in which uncertainty can be an important determinant of behavior: miniature (palm-sized) legged robots, and small-scale quadrotor aerial vehicles. Our motivation for using miniature legged robots is based on the stringent size and weight specifications which impose constraints on the power density of the actuators, thereby prohibiting extensive reliance on feedback control to minimize the effect of noise (Hoover et al, 2010). The second class of systems we consider in this work involves small-scale quadrotors. These robots are very versatile platforms, and their potential in real-world applications has been demonstrated in many recent publications (Gillula et al, 2011; Kumar and Michael, 2012; Lupashin et al, 2010; Mellinger et al, 2012). One source of uncertainty that affects the motion of such systems is due to aerodynamics (Powers et al, 2012; Zarovy et al, 2010), which are difficult to model and capture in low-dimensional models, such as those typically used in feedback control (Kumar and Michael, 2012) and motion planning (Pivtoraiko et al, 2013).

With the method reported here, one can systematically extend an underlying deterministic model to a stochastic regime, and validate the outcome of this procedure against experimen-

The authors are with the Department of Mechanical Engineering, University of Delaware. Email: {kkaryd, poulakas, jxsun, btanner}@udel.edu

tal data. In particular, given a model and experimental data, the method provides a way to estimate the magnitude of the uncertainty that needs to be infused in the model in order to capture the range of behaviors observed in experiments, while providing probabilistic guarantees on the validity of the reported model output.

Such probabilistic guarantees of model performance are important for making informed choices on the type of appropriate planners and controllers (Censi et al, 2008). In addition, once the data variability is captured within the model, analytic tools can be brought to bear for propagating the uncertainty as the physical process evolves (Chirikjian, 2012). This can be useful in application such as motion planning in the presence of uncertainty (van den Berg et al, 2011; Blackmore et al, 2011; Kewani et al, 2009; Lupashin et al, 2010; Pivtoraiko et al, 2013; Thrun et al, 2005), and filtering and estimation (Abbeel et al, 2005; Long et al, 2012; Thrun et al, 2005).

Furthermore, the approach described here can support applications in which data variability plays a significant role, but is either assumed or provided without any probabilistic guarantees. In essence, the method provides a constructive way to tune the variability of the model's output, based on experimental data, at given probabilistic guarantees on model fidelity. Robotics applications in this realm that can benefit from probabilistically validated models include stochastic control design (Anderson and Milutinović, 2014; Shah et al, 2012); Linear Temporal Logic (LTL) control of uncertain vehicles (Cizelj and Belta, 2014); model verification (Steinhardt and Tedrake, 2012) and calibration (Seegmiller et al, 2013); and needle steering (Alterovitz et al, 2008). Moreover, the method here can supplement stochastic trajectory optimization techniques like STOMP (Kalakrishnan et al, 2011), by identifying—based on data—uncertain terms in the optimization function used therein.

The rest of the paper is organized as follows. Section II reports on related work, and Section III presents the proposed framework for simultaneous stochastic extension and probabilistic validation. Then, the applicability of the methodology is demonstrated in two profoundly different case studies: Section IV deals with a miniature legged robot modeled by a stride-to-stride stochastic kinematic map, whereas Section V analyzes flight stability aspects in relation to a small-scale quadrotor, modeled by a stochastically perturbed differential equation. Section VI concludes.

II. Related Work

System identification techniques focus on learning models and fitting parameters to available data, and offer bounds on the fitting and out-of-sample generalization errors. For instance, linear system identification approaches assign weights to available data and identify their optimal values for linear classification, and linear and logistic regression (Abu-Mostafa et al, 2012). If state-space models are required, Linear Time Invariant (LTI) system models can be also obtained (Ljung, 1999). The use of linear models as building blocks supports more powerful nonlinear formulations. For example, cascade products of linear models can generate neural networks (Haykin, 1999), and suitable nonlinear transformations give rise to kernel methods (Hofmann et al, 2008), such as Volterra models (Ogunfunmi, 2007; Schetzen, 2006). Genetic algorithms can distill physical laws by selecting nonlinear terms in ODE models (Schmidt and Lipson, 2009); see also Murphy (2012) for a general overview. However, the models produced typically treat uncertainty as noise, which is either filtered out completely or is used to construct worst-case error bounds.

Model validation, for example (Prajna, 2006), uses experimental data, a model of the uncertainty, and a nominal model with its associated error bounds generated by system identification, to report on whether the proposed model can be trusted. These techniques result in *hard* model (in)validation, in the sense that they provide a yes or no answer to the question of whether a model captures the available data. However, these methods do not provide sufficient insight on the frequency of the events that result in model invalidation; having this information can be useful for refining the model. Hard model (in)validation can be relaxed in a probabilistic sense by employing tools from statistical learning theory (Mendelson, 2003; Vapnik, 1998). Some applications involve correlation analysis of residuals (Ljung, 1999), prediction error within a robust control framework (Gevers et al, 2003), and computation of relative weighted volumes of convex sets for parametric uncertainty models (Lee and Poolla, 1996). A different approach employs a probabilistic model validation methodology to compare a model-generated output probability density function (PDF) with one observed through experiments (Halder and Bhattacharya, 2014). The approach relies on the availability of analytic expressions for propagating the uncertainty through the model at hand, and provides sample-complexity bounds for robust validation inference based on randomized algorithms.

Randomized algorithms offer computationally tractable means to tackle problems in control synthesis (Calafiore et al, 2011; Koltchinskii et al, 2000; Vidyasagar, 2001, 2003), neural networks (Vidyasagar, 2003), and robustness analysis (Alamo et al, 2009; Ray and Stengel, 1993). Typically, deriving worst-case (robust) bounds usually requires a large body of experimental data for theoretical guarantees to hold (Vidyasagar, 2001). However, it has been found that these bounds can be relaxed at the expense of introducing a probabilistic risk, captured by the notion of the *probability of violation* (Alamo et al, 2009; Calafiore et al, 2011; Tempo et al, 2012). This concept can be used to allow some design specifications to be violated, albeit with a relatively small probability. In this way, the sample complexity significantly decreases, at the cost of accepting a risk of violation. This idea has been used in system identification to optimally discard sets of small measure from the set of deterministic estimates of design parameters in Dabbene et al (2012a,b).

In our work, we employ the notion of probability of violation to turn deterministic models into augmented stochastic models, validating that the latter capture the variability observed in experimental data. The approach developed here involves a set-membership characterization of the output PDFs, and it applies directly to a wide range of target models, irrespectively of whether they are phenomenological or derived

based on first principles. We demonstrate this by applying the framework on two distinct problems: (i) to estimate the parameters of an abstract stochastic kinematic model for a miniature legged robot, and (ii) to capture the ground aerodynamic effects in quadrotor ODE models.

The work in this paper is conceptually related to that of Hall et al (2012), where an underlying model is used to provide prior information when training a target Gaussian Process model (Rasmussen and Williams, 2006) based on the efficient PILCO algorithm (Deisenroth and Rasmussen, 2011). However, the predictive ability of that target model deteriorates significantly when the operating point is shifted even slightly and enters to an area where no data are available. As we show in Section V, this paper's method is more robust in the sense that the resulting stochastic extension can make accurate predictions in different operating points, provided that the induced operating conditions do not change the nature of the mechanisms by which uncertainty affects the system; this is attributed to that our approach makes direct use of a deterministic model that relates to the physics of the underlying process.

III. Joint Stochastic Model Extension and Probabilistic Validation

The main ingredients of the proposed framework are described in this section. A general account of the method is first presented in Sections III-A through III-D. A tractable algorithm is then formulated, and made concrete once some assumptions on the underlying statistical distributions are made, in Section III-E. Interspersed between the stages of the conceptual development, are a number of comments that connect the discussion to the examples of Sections IV and V.

A. Overview

Consider a sample space $\mathcal{W}$ that includes all possible outcomes generated by experiments, where observations are collected from a dynamical process of interest. Each element $w \in \mathcal{W}$ consists of state observations obtained during a single experiment. For example, in Section IV $w \in \mathcal{W}$ will be a motion path for the geometric center of the miniature legged robot OctoRoach (Pullin et al, 2012), when it is implementing a specific open-loop controller.

Suppose that a model $\mathcal{M}$ is available for the dynamical process of interest. The model is parameterized by $p \in \mathbb{N}$ parameters, which are collected in a vector ξ taking values in $\Xi \subset \mathbb{R}^p$. For example, in Section IV again, $\mathcal{M}$ takes the form of a stride-to-stride map, while in Section V it is a set of differential equations modeling vertical quadrotor flight. As $\xi \in \Xi$ varies, a family of models $\{\mathcal{M}(\xi),\ \xi \in \Xi\}$ is generated; we will refer to each member $\mathcal{M}(\xi)$ obtained for a specific $\xi \in \Xi$ as a *model instantiation* and we will denote $\mathsf{out}(\mathcal{M}(\xi))$ its output.

Typically, given a collection of $I \in \mathbb{N}$ samples $\{w_1, ..., w_I\}$ obtained experimentally, where each $w_i \in \mathcal{W}$, one can compute the value $\bar{\xi} \in \Xi$ of the model parameters that results in a model instantiation $\mathcal{M}(\bar{\xi})$, the output $\mathsf{out}(\mathcal{M}(\bar{\xi}))$ of which best reproduces the average of the experimentally observed system behavior. One way to find $\bar{\xi}$ is by solving a least-squares problem

$$\bar{\xi} = \arg\min_{\xi \in \Xi} \sum_{t=1}^{T} \|\mathsf{out}(\mathcal{M}(\xi))_t - w^{\mathrm{ave}}(t)\|^2 \ , \qquad (1)$$

where $||\cdot||$ denotes the Euclidean norm, w^{ave} is the average of the set of samples w_i, for $i \in \{1, .., I\}$, and $t \in 1, \ldots, T$ is used here to emphasize that both the output of the model instantiation, and the experimental data, are expressed in time series form of length T. Later in Section III-D, this time dependence is made even more explicit.

In many applications, knowing merely the value $\bar{\xi}$ of the parameter vector that results in a best-fit model instantiation $\mathcal{M}(\bar{\xi})$, may not be sufficient. For example, when using a model to plan the motion of a robot in the presence of uncertainty—as is done for instance by van den Berg et al (2011) and Pivtoraiko et al (2013)—one needs to know not only the average path behavior, but also the paths' distribution around this average. Only then can one quantify the probability that the robot collides with obstacles.

The purpose of this work is to provide a new tool that extends deterministic models to a stochastic regime based on experimental data, and provides probabilistic guarantees of validity in doing so. The resulting stochastic model is considered valid when (i) it has low probabilistic risk of producing a response that is not consistent with the experimental data (model fidelity), and (ii) the resulting paths cover as much as possible of the area marked by the experimental data (model expressiveness). Consequently, this procedure inform us about the range of nominal model parameters, *and* the uncertainty that needs to be infused in the model, so that it can jointly reproduce the experimental data on average and capture the observed variability. The method proposed in this work can be applied to a wide range of models of physical processes, not necessarily expressed in the form of differential equations.

B. Quantifying Model Fidelity: The Probability of Violation

Associated with the sample space $\mathcal{W}$ is a probability measure $\mathbb{P}_{\mathcal{W}}$, which reflects one's belief regarding how the data of the physical system are distributed in $\mathcal{W}$. For example, in the case of the OctoRoach implementing an open-loop straight-line motion controller, one expects the majority of the experimentally produced paths to be clustered around a straight line, and hence $\mathbb{P}_{\mathcal{W}}$ should "peak" on this line. It should be emphasized though, that the proposed method does not depend on the specific form of $\mathbb{P}_{\mathcal{W}}$—which is dictated by the physics of the problem—and it can be applied irrespectively of how $\mathbb{P}_{\mathcal{W}}$ is approximated.

Given $\mathcal{W}$ and $\mathbb{P}_{\mathcal{W}}$, a *multisample* $\mathbf{w}$ is defined as a collection of $K \in \mathbb{N}$ independent and identically distributed (i.i.d.) samples w_k, $k \in \{1, ..., K\}$, drawn from $\mathcal{W}$ according to $\mathbb{P}_{\mathcal{W}}$, and is denoted $\mathbf{w} = \{w_1, ..., w_K\}$. Thus, the multisample $\mathbf{w}$ is drawn from the Cartesian product $\mathcal{W}^K = \mathcal{W} \times \cdots \times \mathcal{W}$ (K-times) according to the probability measure $\mathbb{P}_{\mathcal{W}^K}$. In our setting, we generate multisamples by repeating an experiment

K times, assuming that each experiment is independent[1] of others and that all experiments are performed under identical conditions. To provide some intuition, Fig. 1(a) highlights a multisample of $K = 8$ sample paths obtained by implementing a straight-line controller on the miniature legged robot OctoRoach; see Section IV.

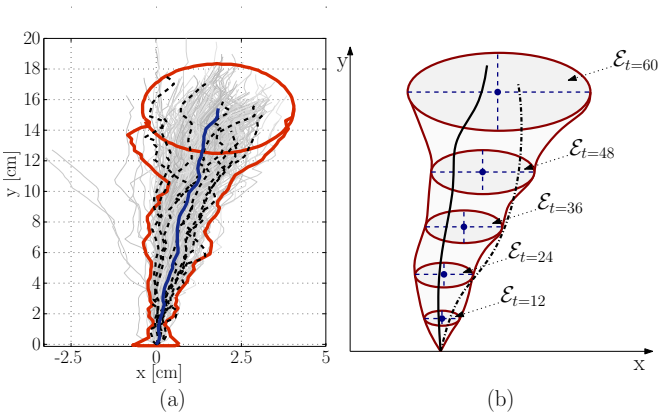

Figure 1. (a) A multisample of length $K = 8$. The sample paths of interest are marked with dashed curves, and are superimposed on top of the whole experimental set of paths. The thick solid curve in the center denotes the average of the eight sample paths, while the thick outline denotes the corresponding *cone of data*, explained below in Section III-D1. (b) A schematic representation for computing the cone of data and the decision function. For each $t \in \{1, \ldots, 60\}$, the *data variability ellipses* $\mathcal{E}_t$ are centered at the sample mean (marked with disks), while their axes are constructed based on sample variances of the multisample; see Sections III-D1 and III-E. Taking the union of all the ellipses yields the cone of data for a particular multisample. Then, the decision function g reports 0 if a model instantiation never crosses the boundary of the cone of data, as shown with the thick curve, and 1 ("violation") otherwise (as shown with the dashed curve crossing the boundary at $t = 24$).

Given a value $\xi \in \Xi$ for the model parameters, we are interested in making a decision as to whether the corresponding model instantiation $\mathcal{M}(\xi)$ will be in agreement, at any time, with the experimentally obtained data. To achieve this, we define a binary-valued decision function $g : \mathcal{W}^K \times \Xi \to \{0, 1\}$ that effectively measures the extent to which the output $\mathsf{out}(\mathcal{M}(\xi))$ of the model instantiation $\mathcal{M}(\xi)$ computed for particular $\xi \in \Xi$ is representative of the data that form the multisample $\mathbf{w} \in \mathcal{W}^K$. To make this statement more precise, we say that the model's output $\mathsf{out}(\mathcal{M}(\xi))$ for $\xi \in \Xi$ is representative of the data forming a multisample $\mathbf{w} \in \mathcal{W}^K$, when this output falls within a prespecified confidence region at level $\gamma \in (0, 1)$. The confidence region is evaluated based on the data in $\mathbf{w}$ and is centered around the multisample's mean. The area covered by the confidence region is called the *cone of data* and is denoted $\mathsf{cone}_\gamma(\mathbf{w})$. Section III-D1 below provides details on computing $\mathsf{cone}_\gamma(\mathbf{w})$ in a general setting, while Section III-E shows how these computations can become tractable by particularizing on the measure $\mathbb{P}_{\mathcal{W}^K}$ (see Figure 1(b)).

[1] As is often the case, the assumption of independence is difficult to justify in practice. Note though that certain properties that are relevant to our discussion can be extended when the sequence of samples is not i.i.d. but satisfies a "mixing" condition (Vidyasagar, 2003, Section 2.5, p. 33). We will not discuss this issue further, for it requires the introduction of a number of technical results that would shift the focus of this work; the interested reader is referred to Vidyasagar (2003, Chapter 3).

The decision function g can now be defined as

$$g(\mathbf{w}, \xi) := \begin{cases} 0, & \text{if } \mathsf{out}(\mathcal{M}(\xi)) \subset \mathsf{cone}_\gamma(\mathbf{w}) \\ 1, & \text{otherwise} \end{cases} . \tag{2}$$

Intuitively, the function g is interpreted as a penalty on a model instantiation $\mathcal{M}(\xi)$ for $\xi \in \Xi$ whenever $\mathcal{M}(\xi)$ produces a behavior that is statistically different from experimental data $\mathbf{w}$. For a given $\xi \in \Xi$, consider the set

$$A_\xi := \{\mathbf{w} \in \mathcal{W}^K \mid g(\mathbf{w}, \xi) = 1\} \tag{3}$$

which contains all the multisamples that violate the condition $\mathsf{out}(\mathcal{M}(\xi)) \subset \mathsf{cone}_\gamma(\mathbf{w})$. As $\xi \in \Xi$ varies, a collection of sets

$$\mathcal{A} := \{A_\xi, \ \xi \in \Xi\}$$

in $\mathcal{W}^K$ is generated, each of which contains the "bad" multisamples for the corresponding parameter values ξ. The *probability of violation* can then be defined by the function $P : \Xi \to [0, 1]$ given by the rule

$$P(\xi) := \mathbb{P}_{\mathcal{W}^K}(A_\xi) \ , \tag{4}$$

which provides a measure of the subset $A_\xi \subset \mathcal{W}^K$ of multisamples that are statistically inconsistent with the particular model instantiation $\mathcal{M}(\xi)$. More precisely, for a *given* value $\xi \in \Xi$ of the model's parameters, $P(\xi)$ expresses the likelihood of generating a multisample $\mathbf{w}$ by sampling $\mathcal{W}^K$ according to $\mathbb{P}_{\mathcal{W}^K}$, which—for a desired confidence level $\gamma \in (0, 1)$—results in a $\mathsf{cone}_\gamma(\mathbf{w})$ that does not include the output $\mathsf{out}(\mathcal{M}(\xi))$ of the model instantiation $\mathcal{M}(\xi)$.

The probability of violation (4) provides a means of deciding whether a *specific* model instantiation $\mathcal{M}(\xi)$ is probabilistically consistent with experimental data. For a given ξ, a large value for $P(\xi)$ implies that the chance of generating multisamples that are not in agreement with the particular model instantiation $\mathcal{M}(\xi)$ is high, suggesting that the fidelity of $\mathcal{M}(\xi)$ is low.

In addition to quantifying model fidelity, we also want to express multisample variability; we cannot do this with a single model instantiation—we need to look at distributions of model instantiations over model parameters.

C. Distributions of Models and Model Expressiveness

One way in which a model can capture the dispersion of the experimentally generated data is through stochasticity in the model's parameters: their value is the outcome of a random experiment. In particular, consider a sample space Ω_Ξ containing all the possible outcomes of such random experiment and a probability measure $\mathbb{P}_\Xi$ belonging in a family of measures $\mathcal{P}_\Xi$. Then, the parameters of the model form a random vector $\tilde{\xi} : \Omega_\Xi \to \Xi$, the realization of which results in the values $\xi \in \Xi$ that determine the model instantiation $\mathcal{M}(\xi)$. With this construction, for each $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$ a *distribution* of model instantiations $\mathcal{D}_{\mathbb{P}_\Xi} = \{\mathcal{M}, \Xi, \mathbb{P}_\Xi\}$ is defined.

Note that each measure $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$ is assumed here to "peak" around the parameter value $\bar{\xi}$ obtained by the solution of the least-squares optimization problem (1). This reflects our intuition that the output of $\mathcal{M}(\bar{\xi})$ is a good representation of the average of the experimental measurements. One can

choose the family of measures $\mathcal{P}_\Xi$ in a way that reflects their own beliefs on how the stochasticity enters the nominal model parameters. Also, the dispersion of the values ξ obtained by sampling the random vector $\tilde{\xi}$ according to different measures in $\mathcal{P}_\Xi$ can be different. As a result, the variability of the outputs produced by the model instantiations $\mathcal{M}(\xi)$ generated by sampling a distribution $\mathcal{D}_{\mathbb{P}_\Xi}$ varies across the collection $\{\mathcal{D}_{\mathbb{P}_\Xi},\ \mathbb{P}_\Xi \in \mathcal{P}_\Xi\}$. The problem now reduces to identifying the distribution $\bar{\mathcal{D}}_{\mathbb{P}_\Xi} \in \{\mathcal{D}_{\mathbb{P}_\Xi},\ \mathbb{P}_\Xi \in \mathcal{P}_\Xi\}$ which captures best the variability in the experimental data without violating a desired specification on the probability of violation of the model instantiations contained in $\bar{\mathcal{D}}_{\mathbb{P}_\Xi}$. The purpose of this section is to make this statement precise.

We begin by providing a way to evaluate the capacity of a distribution of models $\mathcal{D}_{\mathbb{P}_\Xi}$ for capturing the experimental data, given $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$. Consider the set

$$S := \{\xi \in \Xi \mid P(\xi) \geq P_0\} \ , \tag{5}$$

which includes the parameter values ξ that result in model instantiations $\mathcal{M}(\xi)$, each corresponding to a probability of violation exceeding $P_0 \in (0,1)$. Note that the size of the set S depends on the measure $\mathbb{P}_\Xi$, implying that the likelihood of model instantiations that satisfy (5) is different for different distributions $\mathcal{D}_{\mathbb{P}_\Xi}$. Given a desired level $\alpha \in [0,1)$ we require

$$\mathbb{P}_\Xi(S) \leq \alpha \ . \tag{6}$$

Clearly, a lower value of the parameter α corresponds to stricter fidelity standards for the model instantiations $\mathcal{M}(\xi)$ generated by sampling $\mathcal{D}_{\mathbb{P}_\Xi}$. In fact, selecting $\alpha = 0$ implies $P_0 = \sup_{\xi\in\Xi} P(\xi)$, since the probability that a model instantiation $\mathcal{M}(\xi) \in \mathcal{D}_{\mathbb{P}_\Xi}$ results in a probability of violation that exceeds P_0 is required to be zero. Hence, $\alpha = 0$ corresponds to the most conservative way of tuning the behavior expressiveness of a distribution of models $\mathcal{D}_{\mathbb{P}_\Xi}$; that is, $\mathcal{D}_{\mathbb{P}_\Xi}$ is characterized by the model instantiation with the worst performance, in the sense that the corresponding parameters maximize the probability of violation.

In view of (5), relaxing α in (6) means that we allow a set $S \subset \Xi$ of parameter values to be exceptions to the fidelity rule. The size of S can be explicitly controlled through (6) by selecting α so that (Vidyasagar, 2001, Section 3)

$$\sup_{\xi\in\Xi\setminus S} P(\xi) \leq P_0 \leq \sup_{\xi\in\Xi} P(\xi) \ ; \tag{7}$$

i.e., P_0 is bounded from above by the supremum of the probability of violation over all parameter values (most stringent characterization) and from below by the supremum of the probability of violation over "nearly" all parameter values (turning a blind eye to parameters in S). The implication of (7) is that P_0 is a *probable near maximum* (Type 2 near maximum) over the set Ξ of the probability of violation $P(\cdot)$ to the level α (Vidyasagar, 2003, Definition 11.2, p. 433); (Vidyasagar, 2001, Section 3). Note that P_0 depends *both* on the measure $\mathbb{P}_\Xi$ and on the level α; writing $P_0(\mathbb{P}_\Xi, \alpha)$ emphasizes this dependence.

Remark 1: While the value of the probability of violation (4) defined in Section III-B for a given $\xi \in \Xi$ quantifies the behavior of a single model instantiation $\mathcal{M}(\xi)$, the probable near maximum P_0 to the level α of $P(\cdot)$ over a distribution of models $\mathcal{D}_{\mathbb{P}_\Xi}$ provides a measure of how faithfully $\mathcal{D}_{\mathbb{P}_\Xi}$ captures the experimental data. Indeed, for a given level α, the smaller P_0 is for a distribution $\mathcal{D}_{\mathbb{P}_\Xi}$, the more faithful this distribution is in capturing the data.

We are now ready to provide a precise formulation of the problem described at the beginning of the section: Given (i) a family of model distributions $\{\mathcal{D}_{\mathbb{P}_\Xi},\ \mathbb{P}_\Xi \in \mathcal{P}_\Xi\}$, (ii) a level $\alpha \in [0,1)$, and (iii) a desired fidelity specification $\rho \in [0,1)$, determine the distribution $\bar{\mathcal{D}}_{\mathbb{P}_\Xi}$—or, equivalently, the corresponding measure $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$—that maximizes the dispersion of a random vector $\tilde{\xi}$ of model parameters, provided that the probable near maximum to level α of the probability of violation does not exceed ρ. Mathematically, this translates to finding the measure $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$ that realizes

$$\sup_{\mathbb{P}_\Xi\in\mathcal{P}_\Xi} \mathrm{Tr}\left(\mathrm{Cov}(\tilde{\xi},\tilde{\xi})\right) \tag{8}$$

subject to the constraint

$$P_0(\mathbb{P}_\Xi, \alpha) \leq \rho \ , \tag{9}$$

where $\mathrm{Cov}(\tilde{\xi},\tilde{\xi})$ is the covariance matrix associated with the random vector $\tilde{\xi}$ and $\mathrm{Tr}(\cdot)$ denotes trace.

Implicit here is the assumption that as the measure $\mathbb{P}_\Xi$ changes to make the variance on $\tilde{\xi}$ grow, the model instantiations generated by sampling $\tilde{\xi}$ produce outputs $\mathsf{out}(\mathcal{M}(\xi))$ that are more and more dispersed. Hence, the solution of (8)–(9) is expected to result in a distribution over the model parameters ξ that allows the corresponding distribution of model instantiations $\mathcal{D}_{\mathbb{P}_\Xi}$ to reproduce, at a given confidence level, as many experimental behaviors observed as possible—not only the average.

D. A Randomized Approach for Stochastic Model Extension and Probabilistic Validation

This section provides details on computing the decision function g defined by (2) in a general context, and proposes a randomized approach for estimating the quantities involved in the implementation of the method via explicit computations.

1) Cone of Data and Decision Function: Consider a multisample $\mathbf{w} = \{w_1, ...w_K\} \in \mathcal{W}^K$ generated experimentally by executing an experiment K times. Each $w_k \in \mathbf{w}$ has the form of a time series

$$w_k = \{(x_{k,1}(t), x_{k,2}(t), ..., x_{k,L}(t))\}_{t\in\{1,...,T\}} \ ,$$

where $x_{k,\ell}(t)$ is a measurement at time $t \in \{1, ..., T\}$, of the system state indexed $\ell \in \{1, .., L\}$, during the experiment $k \in \{1, ..., K\}$. Time instants $t \in \{1, ..., T\}$ are determined based on the sampling frequency of data collection.

We can associate to each multisample, $\mathbf{w}$, a confidence interval $\mathcal{I}_{\ell,t}(\mathbf{w})$ to which an experimental trial belongs according to a given probability. A typical way to construct such intervals is to use information about the underlying distribution; for example, in Section III-E below we assume a normal distribution and use the sample mean and variance to construct confidence intervals. Another way is based on bootstrap (Efron and Tibshirani, 1994). This class of methods

relies on resampling of the original data, and can be used to estimate sample distributions of various statistics.

Then, the intervals $\mathcal{I}_{\ell,t}(\mathbf{w})$ provide the basis for constructing the L-dimensional *data variability ellipsoid* $\mathcal{E}_t(\mathbf{w})$. At each time instant $t \in \{1,...,T\}$, the ellipsoid $\mathcal{E}_t(\mathbf{w})$ is centered at the point $(\bar{x}_1(t), \bar{x}_2(t), ..., \bar{x}_L(t))$, and $\mathcal{I}_{\ell,t}(\mathbf{w})$ for $\ell \in \{1,\dots,L\}$ are its principal axes (see Fig. 1(b)). Note that the dependence of $\mathcal{E}_t$ on $\mathbf{w}$ appears explicitly to highlight the fact that these constructions are specific to a given multisample. The cone of data corresponding to a multisample $\mathbf{w} \in \mathcal{W}^K$ at level $\gamma \in (0,1)$ is then the union of all L-dimensional ellipsoids

$$\mathsf{cone}_\gamma(\mathbf{w}) = \bigcup_{t=1}^{T} \mathcal{E}_t(\mathbf{w}) \ . \tag{10}$$

Figure 1(b) provides a schematic representation of the cone of data associated with a multisample, for $T = 60$, for the case of the OctoRoach robot studied in Section IV.

To evaluate the decision function g defined by (2) for a $\xi \in \Xi$ given $\mathbf{w}$, we need to specify how we check the condition $\mathsf{out}(\mathcal{M}(\xi)) \subset \mathsf{cone}_\gamma(\mathbf{w})$ for the associated model instantiation $\mathcal{M}(\xi)$. We work element-wise through $\{1,\dots,T\}$, first by setting

$$\mathsf{out}(\mathcal{M}(\xi)) = \{\mathsf{out}(\mathcal{M}(\xi))_t\}_{t\in\{1,...,T\}} \ ,$$

where for any $t \in \{1,\dots,T\}$

$$\mathsf{out}(\mathcal{M}(\xi))_t := (x_{\mathcal{M},1}(t), \dots, x_{\mathcal{M},L}(t)) \ ,$$

and then by defining the indicator function

$$\mathbf{1}_{\mathcal{E}_t(\mathbf{w})}(\mathsf{out}(\mathcal{M}(\xi))_t) := \begin{cases} 1, & \text{if } \mathsf{out}(\mathcal{M}(\xi))_t \in \mathcal{E}_t(\mathbf{w}) \\ 0, & \text{otherwise} \end{cases} \tag{11}$$

that checks the inclusion condition at every $t \in \{1,\dots,T\}$. In this way the decision function is found as

$$g(\mathbf{w},\xi) = 1 - \prod_{t=1}^{T} \mathbf{1}_{\mathcal{E}_t(\mathbf{w})}(\mathsf{out}(\mathcal{M}(\xi))_t) \ . \tag{12}$$

Note that (12) requires the inclusion to hold for all time instants $t \in \{1,...,T\}$; if at any single t $\mathsf{out}(\mathcal{M}(\xi))_t \notin \mathcal{E}_t(\mathbf{w})$, the decision function is triggered and the model is considered to have violated the fidelity specification.

2) Approximating the Probability of Violation: The probability of violation (4) is difficult to compute explicitly, even if the probability measure $\mathbb{P}_{\mathcal{W}^K}$ is analytically available. However, this probability can be effectively approximated empirically (Alamo et al, 2009). If $\mathbf{W}_M = \{\mathbf{w}_1, ..., \mathbf{w}_M\} \in (\mathcal{W}^K)^M$ is a collection of M multisamples of length K, each drawn from $\mathcal{W}^K$, the empirical probability of violation is

$$\hat{P}(\xi; \mathbf{W}_M) = \frac{1}{M} \sum_{m=1}^{M} g(\xi, \mathbf{w}_m) \ , \tag{13}$$

where the dependence of $\hat{P}$ on both the specific collection of multisamples $\mathbf{W}_M = \{\mathbf{w}_1, ..., \mathbf{w}_M\} \in (\mathcal{W}^K)^M$ and on the parameter values $\xi \in \Xi$ that determine the violation set $A_\xi \in \mathcal{A}$ in (3) appears explicitly. Note that $\hat{P}(\xi; \mathbf{W}_M)$ is a random variable. For $\varepsilon > 0$, consider (Vidyasagar, 2001, Section 4)

$$q(M, \varepsilon, \mathbb{P}_{\mathcal{W}^K}) := \mathbb{P}_{(\mathcal{W}^K)^M}\big\{\mathbf{W}_M \in (\mathcal{W}^K)^M : \sup_{\xi\in\Xi} |\hat{P}(\xi; \mathbf{W}_M) - P(\xi)| > \varepsilon\big\} \ . \tag{14}$$

Then, $1 - q(M, \varepsilon, \mathbb{P}_{\mathcal{W}^K})$ is the confidence with which we can say that $\hat{P}(\xi, \mathbf{W}_M)$ is within $\varepsilon > 0$ of the true $P(\xi)$. If $q(M, \varepsilon, \mathbb{P}_{\mathcal{W}^K}) \to 0$ as $M \to \infty$ for any fixed ε, then the empirical probabilities converge uniformly to their true values, implying that the collection of sets $\mathcal{A}$ has the property of uniform convergence of empirical probabilities (UCEP); see (Vidyasagar, 2003, Section 3.1, p. 45). Establishing the UCEP property for the collection of sets $\mathcal{A}$ can be difficult if $\mathcal{A}$ is infinite; but if this is a finite collection, Hoeffding's inequality (Vidyasagar, 2003, Lemma 2.7, p. 26) yields

$$q(M, \varepsilon, \mathbb{P}_{\mathcal{W}^K}) \leq 2|\mathcal{A}| \exp\left(-2M\varepsilon^2\right) \ , \tag{15}$$

where $|\mathcal{A}|$ is the cardinality of $\mathcal{A}$. In fact, since $\{2|\mathcal{A}| \exp(-2M\varepsilon^2)\}_{M\in\mathbb{N}}$ is summable (Vidyasagar, 2003, Lemma 2.10, p. 31), $\hat{P}(\xi; \mathbf{W}_M)$ not only converges uniformly to $P(\xi)$ with M, but also almost surely. The inequality (15) can be used to provide bounds for the sample size M that achieves the desired accuracy and confidence specifications.

3) Approximating the Maximum of the Probability of Violation over a Distribution of Models: We have seen in Section III-C that the expressiveness of a distribution of model instantiations $\mathcal{D}_{\mathbb{P}_\Xi} = \{\mathcal{M}, \Xi, \mathbb{P}_\Xi\}$ for a given probability measure $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$ can be characterized by evaluating a probable near maximum P_0 at level α, of the probability of violation $P(\cdot)$ over $\mathcal{D}_{\mathbb{P}_\Xi}$.

With the probabilistic setting of Section III-C, a collection of $N \in \mathbb{N}$ samples ξ_n, $n \in \{1, ..., N\}$ is drawn according to $\mathbb{P}_\Xi$, thereby resulting in a *parameter multisample* denoted by $\boldsymbol{\xi}_N = \{\xi_1, ..., \xi_N\}$. Note that the parameter multisample $\boldsymbol{\xi}_N$ is drawn from the Cartesian product $\Xi^N = \Xi \times \cdots \times \Xi$ (N-times) according to the measure $\mathbb{P}_{\Xi^N}$. Then, for a given (fixed) *data multisample* $\mathbf{w} \in \mathcal{W}^K$, the probability of violation is a random variable due to its dependence on $\tilde{\xi}$, and $\{P(\xi_1), ..., P(\xi_N)\}$ are the corresponding samples of the probability of violation. Now define

$$\bar{P}_0(\boldsymbol{\xi}_N) := \max_{n\in\{1,..,N\}} P(\xi_n) \ , \tag{16}$$

and consider the set

$$\bar{S}_{\boldsymbol{\xi}_N} := \big\{\xi \in \Xi \mid P(\xi) > \bar{P}_0(\boldsymbol{\xi}_N)\big\} \ , \tag{17}$$

which is defined similarly to S in (5), only now we have used P_0 instead of $\bar{P}_0(\boldsymbol{\xi}_N)$, computed on the basis of the parameter multisample $\boldsymbol{\xi}_N$. Then, Vidyasagar (2003, Lemma 11.1, p. 427) asserts that

$$\mathbb{P}_{\Xi^N}\big\{\boldsymbol{\xi}_N \in \Xi^N \mid \mathbb{P}_\Xi(\bar{S}_{\boldsymbol{\xi}_N}) > \alpha\big\} \leq (1-\alpha)^N \ . \tag{18}$$

The inequality implies that $\bar{P}_0(\boldsymbol{\xi}_N)$ is an "empirical estimate" of the supremum of the probability of violation $P(\cdot)$ over $\mathcal{D}_{\mathbb{P}_\Xi}$. However, $\bar{P}_0(\boldsymbol{\xi}_N)$ is a different type of estimate compared to $\hat{P}(\xi; \mathbf{W}_M)$ in (13), because (18) does not require $\bar{P}_0(\boldsymbol{\xi}_N)$ to converge uniformly to the true supremum of $P(\cdot)$. Rather, the

claim is that the probability that the violation set $\bar{S}_{\boldsymbol{\xi}_N}$ has small measure, is high. Inequality (18) is used to specify the size N of the parameter multisample $\boldsymbol{\xi}_N$.

E. Algorithm and Implementation

The results described in the previous sections did not assume any particular type of probability distribution. To facilitate computations, however, in what follows, we will assume that $\mathbb{P}_{\mathcal{W}^K}$ corresponds to a (joint) Gaussian distribution.[2] This assumption greatly simplifies the computation of the cone of data associated with a specific multisample $\mathbf{w} \in \mathcal{W}^K$.

Under a Gaussian assumption, each multisample of size K can be used to estimate the average of the population of data within a given confidence level $\gamma \in (0,1)$. For each $\ell \in \{1, ..., L\}$, let

$$\bar{x}_\ell(t) = \frac{1}{K}\sum_{k=1}^{K} x_{k,\ell}(t) \ ,$$

be the sample average at a given time instant $t \in \{1, ..., T\}$, and denote

$$s_\ell(t) = \sqrt{\frac{1}{K-1}\sum_{k=1}^{K}\left(x_{k,\ell}(t) - \bar{x}_\ell(t)\right)^2}$$

the corresponding sample standard deviation. Then, the confidence interval $\mathcal{I}_{\ell,t}(\mathbf{w})$ at level γ for the ℓ-th variable associated with the multisample $\mathbf{w}$ at the time instant $t \in \{1, ..., T\}$ is

$$\mathcal{I}_{\ell,t}(\mathbf{w}) = [\bar{x}_\ell(t) - \tau_\gamma \frac{s_\ell(t)}{\sqrt{K}},\ \bar{x}_\ell(t) + \tau_\gamma \frac{s_\ell(t)}{\sqrt{K}}] \ , \qquad (19)$$

where the percentile τ_γ at level γ is a constant that can be found in tables for different distributions (Dixon and Massey, 1957).

Table I
NOMENCLATURE

Description	Symbol	Equation
Confidence interval for the cone of data	$\gamma \in (0,1)$	(19)
Level of probable near maximum	$\alpha \in [0,1)$	(6)
Model fidelity specification	$\rho \in (0,1)$	(9)
Accuracy level	$\varepsilon \in (0,1)$	(14)
Confidence level	$\delta \in (0,1)$	(20)-(21)
Number of parameter multisamples	N	(20)
Number of data multisamples	M	(21)
Length of each data multisample	K	–
Total number of experimental paths	$I = M \cdot K$	–
Number of state variables	L	–
Number of model parameters	p	–

To find the appropriate cardinality N of the parameter multisample set ξ_N, we first need to set parameters $\gamma, \alpha, \rho > 0$ and $\varepsilon, \delta > 0$, which we collect in Table I. For computational expediency, each measure $\mathbb{P}_\Xi \in \mathcal{P}_\Xi$ is assumed to be associated with a Gaussian random vector $\tilde{\xi}$ with mean $\mathbb{E}[\tilde{\xi}] = \bar{\xi}$ computed by (1), and covariance matrix $\mathrm{Cov}(\tilde{\xi}, \tilde{\xi}) = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_p^2)$, where p is the number of the parameters. Hence, the family of parameter distributions $\mathcal{P}_\Xi = \{\mathbb{P}_\Xi^{\boldsymbol{\sigma}},\ \boldsymbol{\sigma} \in \mathbb{R}^p\}$ is parameterized by the array $\boldsymbol{\sigma} = \{\sigma_1^2, \ldots, \sigma_p^2\}$ that will be determined by solving the optimization problem (8).

Based on (18), and for $\delta \in (0,1)$, N needs to satisfy

$$(1-\alpha)^N \leq \frac{\delta}{2} \iff N \geq \frac{\log\frac{2}{\delta}}{\log\frac{1}{1-\alpha}} \ , \qquad (20)$$

so that with confidence $1 - \frac{\delta}{2}$, $\bar{P}_0(\boldsymbol{\xi}_N)$ defined by (16) is a probable near maximum of the probability of violation $P(\cdot)$ to a level α. This sampling process results in the finite collection of sets $\mathcal{A} := \{A_{\xi_n},\ \xi_n \in \boldsymbol{\xi}_N\}$ with $|\mathcal{A}| = N$. Then, Hoeffding's inequality (15) links the two sample sizes

$$2N\exp\left(-2M\varepsilon^2\right) < \frac{\delta}{2} \iff M \geq \frac{1}{2\varepsilon^2}\ln\left(\frac{4N}{\delta}\right) \ , \qquad (21)$$

suggesting that if N parameter multisamples are drawn, then M data multisamples need to be obtained experimentally in order for $\hat{P}(\xi; \mathbf{W}_M)$ to be an empirical estimate of $P(\xi)$ with confidence $1 - \frac{\delta}{2}$. Then, we select a measure $\mathbb{P}_\Xi^{\boldsymbol{\sigma}}$ (through $\boldsymbol{\sigma}$) and generate the parameter multisample $\boldsymbol{\xi}_N = \{\xi_1, ..., \xi_N\}$, sampling $\tilde{\xi}$ according to $\mathbb{P}_\Xi^{\boldsymbol{\sigma}}$.

With M data multisamples $\mathbf{w}_m \in \mathbf{W}_M$, and N parameter multisamples $\xi_n \in \boldsymbol{\xi}_N$ available, the decision function (2) is computed explicitly based on (11)-(12). For each ξ_n, $n \in \{1, \ldots, N\}$, (13) results in an empirical probability of violation

$$\hat{P}(\xi_n; \mathbf{W}_M) = \frac{1}{M}\sum_{m=1}^{M} g(\xi_n, \mathbf{w}_m) \ . \qquad (22)$$

Owing to the choice of N and M according to (20) and (21), we can say with confidence $1 - \delta$ that

$$\hat{P}_0 = \max_{n \in \{1,..,N\}} \hat{P}(\xi_n; \mathbf{W}_M) \qquad (23)$$

is a *probably approximate near maximum (Type 3 near maximum)* to accuracy ε and level α (Koltchinskii et al, 2000; Vidyasagar, 2001, 2003) of the probability of violation in (4) over the distribution $\mathcal{D}_{\mathbb{P}_\Xi}$.

Albeit approximate, the procedure described above relaxes in a probabilistic sense the problem of maximizing the function $P(\cdot)$, through an explicitly computable quantity $\hat{P}_0$ that characterizes the expressiveness of a given model distribution $\mathcal{D}_{\mathbb{P}_\Xi}$. The optimization problem defined by (8)–(9) is thus relaxed into the problem of maximizing the variances $\boldsymbol{\sigma} = \{\sigma_1^2, \ldots, \sigma_p^2\}$ in the model parameters within a family of Gaussian distributions, which all "peak" at the solution $\bar{\xi}$ of (1) while satisfying the constraint

$$\hat{P}_0 \leq \rho \ . \qquad (24)$$

The output of the algorithm, summarized in Table II, is the largest $\bar{\boldsymbol{\sigma}} = \{\bar{\sigma}_1^2, \ldots, \bar{\sigma}_p^2\}$ that $\mathcal{M}(\xi)$ can afford before violating (24). Equivalently, the procedure determines the measure $\mathbb{P}_\Xi^{\bar{\boldsymbol{\sigma}}}$ in the form of a Gaussian distribution

[2] We emphasize here that the Gaussianity assumption may not always be the most appropriate choice. In the case of ground mobile robots, for instance, a more appropriate choice for $\mathbb{P}_{\mathcal{W}^K}$ may be the "banana" distribution (Chirikjian, 2012; Long et al, 2012; Thrun et al, 2005)—essentially a Gaussian distribution in exponential coordinates. The Gaussian assumption is imposed here merely for computational expediency; exploring the different options for constructing the cone data is out of the focus of this present paper.

$\mathcal{N}(\bar{\xi}, \mathrm{diag}\{\bar{\sigma}_1^2, \ldots, \bar{\sigma}_p^2\})$, which in turn produces a distribution of models $\mathcal{D}_{\mathbb{P}_\Xi^{\bar{\boldsymbol{\sigma}}}}$ with outputs that cover densely the distribution of experimental data remaining statistically within that cone with confidence γ.

Table II
PROBABILISTICALLY VALID STOCHASTIC MODEL EXTENSIONS ALGORITHM

1.	**Require** $\gamma, \alpha, \rho > 0, \varepsilon, \delta > 0$, and $K \in \mathbb{N}$.
2.	**Calculate** N and M from (20) and (21), respectively.
3.	**Collect** M data multisamples of size K.
4.	**Identify** $\bar{\xi}$ from (1) using all available $I = M \cdot K$ data.
5.	**Select** $\boldsymbol{\sigma}$.
6.	**Generate** $\boldsymbol{\xi}_N = \{\xi_1, \xi_2, \ldots, \xi_N\}$ by sampling $\tilde{\xi} \sim \mathbb{P}_\Xi^{\boldsymbol{\sigma}}$.
7.	**Calculate** $\mathsf{out}(\mathcal{M}(\xi))_t$ for each $\xi_n, n \in \{1, \ldots, N\}$.
8.	**Construct** $\mathsf{cone}_\gamma(\mathbf{w}_m)$ for each $\mathbf{w}_m, m \in \{1, \ldots, M\}$ according to (10)
9.	**Calculate** the empirical probability of violation $\hat{P}(\xi_n; \mathbf{W}_M)$ for each $n \in \{1, \ldots, N\}$ according to (22).
10.	**Select** the probably approximate near maximum $\hat{P}_0$ according to (23).
11.	**If** $\hat{P}_0 \leq \rho$, increase $\boldsymbol{\sigma}$ and **go to** step 6, **else** return $\bar{\boldsymbol{\sigma}}$ and exit.

Remark 2: In the remaining of the paper, we select γ at 95%; this value is a very common choice for the construction of confidence intervals. We also set ρ at 35% in order to produce less conservative results. The values of α, ε, and δ directly determine the amount of experimental trials that should be performed, and different combinations yield the same number of trials; we select α, ε so that the experimental trials are kept at a reasonable number (about 250), and out of all possible combinations, we choose the one that maximizes the confidence $1 - \delta$.

IV. APPLICATION TO MINIATURE LEGGED ROBOTS

As a first example of how the proposed method can be applied, we consider the horizontal-plane stochastic behavior of the miniature legged robot OCTOROACH (Pullin et al, 2012) (Fig. 2), when executing motion plans at low crawling speeds.

Several reduced-order deterministic models—a.k.a. templates—have been studied in the context of robotic legged locomotion, and have been found capable of capturing salient features of the horizontal-plane behavior of miniature legged robots. For instance, the dynamic Lateral Leg Spring (LLS) model (Holmes et al, 2006) has been used to derive turning strategies (Proctor and Holmes, 2008) for the minimally-actuated robot DYNAROACH (Hoover et al, 2010). On another vein, the kinematic Switching Four-bar Mechanism (SFM) model (Karydis et al, 2012, 2015) has been successful in capturing the behavior of the eight-legged robot OCTOROACH (Pullin et al, 2012) when crawling at low speeds, and the Dubins model (Dubins, 1957) has been found capable of predicting the motion of the wheel-legged miniature robot STAR (Zarrouk et al, 2013) at low-sprawled configurations (Karydis et al, 2014).

Unfortunately, these deterministic templates do not capture sufficiently the stochastic nature of the in-plane motion of miniature legged robots. Furthermore, size and weight specifications hinder the ability of the actuators to minimize the effect of noise through feedback control (Hoover et al, 2010). This reduced control authority, combined with the small size and weight of these systems, exacerbate the effect of the inherently stochastic leg-ground interaction, and limit the practical utility of an otherwise satisfactory deterministic model. The reported method provides a way to extend such deterministic templates to a stochastic setting, allowing to capture the inherently stochastic nature of robotic legged locomotion at small scales.

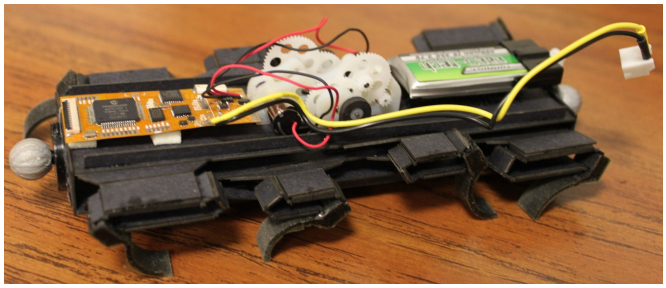

Figure 2. The OCTOROACH, a bio-inspired, eight-legged robot, designed at the University of California, Berkeley. Its length and weight are 130 mm, and 35 g, respectively, and it can reach a maximum speed of 0.5 m/s. The robot features a differential-drive steering method, and employs an alternating tetrapod gait for navigating through its environment.

In this particular example, we focus on the low-speed crawling motion of the OCTOROACH robot (Pullin et al, 2012) (Fig. 2), and apply the proposed framework to stochastically extend the (deterministic) SFM model (Karydis et al, 2012, 2015) by infusing stochasticity in its parameters. Preliminary work focusing only in straight-line motion has been reported in Karydis et al (2013).

A. The Nominal Model

The OCTOROACH is designed to follow an alternating tetrapod gait, as shown in Fig. 3(a). Legs $\{1, 2, 3, 4\}$ form the "right" tetrapod, and legs $\{5, 6, 7, 8\}$ form the "left" tetrapod. The ipsilateral[3] legs of each tetrapod touch the ground at the same instant, and rotate in phase with the same angular velocity, as indicated in the abstract eight-legged model of Fig. 3(b). Collapsing the ipsilateral legs of each tetrapod into a single "virtual" leg, yields the SFM model (Fig. 3(c)) in which the contralateral virtual legs (e.g., $\{O_1, O_2\}$) represent the collective effect of the tetrapod they replace (e.g. $\{1, 2, 3, 4\}$) (Karydis et al, 2012). The SFM model consists of a rigid torso and four rigid legs organized in two pairs, $\{AO_1, BO_2\}$, and $\{AO_3, BO_4\}$ (Fig. 3(c)) called right and left pair, respectively. Assuming no slipping between the legs and the ground, the model follows the footfall pattern of Fig. 4(a), in a gait that couples the torso and the legs into two alternating four-bar linkages.

The model's configuration is $(x_G, y_G, \theta) \in \mathbb{R}^2 \times \mathbb{S}$, where (x_G, y_G) is the position of the geometric center G in some inertial coordinate frame, and θ is the angle the longitudinal body-fixed axis forms with the y-inertial axis. By convention,

[3]Ipsilateral means on the same side and contralateral means on the other side.

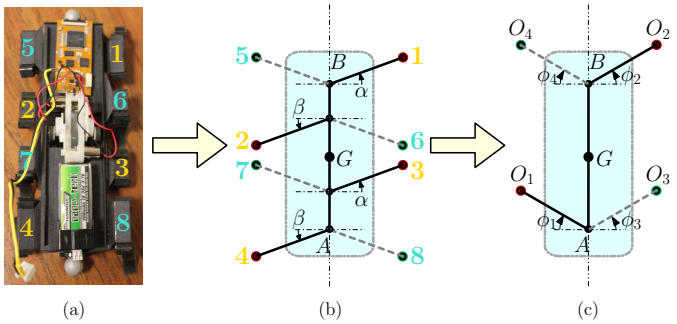

Figure 3. Relating the SFM template to the OctoRoACH. (a) The foot-fall pattern of the robot, is an alternating tetrapod gait. Legs $\{1, 2, 3, 4\}$ form the right tetrapod, and legs $\{5, 6, 7, 8\}$ form the left tetrapod. (b) An eight-legged kinematic simplification of the gait mechanism used by the robot. The ipsilateral legs of each tetrapod are coupled, forming the angles α and β shown here for the right tetrapod. (c) The SFM template is obtained by grouping coupled legs within a tetrapod into a single virtual leg inducing the same displacement. Legs $\{1, 2, 3, 4\}$ reduce to the pair $\{AO_1, BO_2\}$, while legs $\{5, 6, 7, 8\}$ reduce to the pair $\{AO_3, AO_4\}$. The two pairs of legs are activated in turns, thus forming two fourbar linkages, $\{O_1A, AB, BO_2\}$ and $\{O_3A, AB, BO_4\}$.

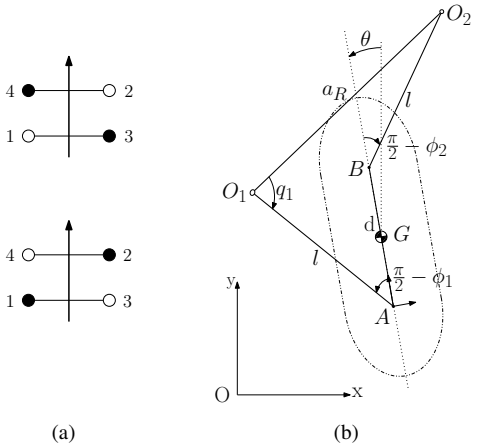

Figure 4. (a) The footfall pattern followed by SFM. (b) Analysis of the model. Due to symmetry, it suffices to analyze only one pair, taken here to be the right. d denotes the distance between the two hip-point joints A and B, and is chosen the same as the length of the robot, that is $d = 13$ cm, l is the leg length, set at 3 cm, G is the model's geometric center, and θ denotes the orientation of the model with respect to some global reference frame.

positive changes in the orientation correspond to counter-clockwise angles θ. During each step, the evolution of this configuration is determined by the kinematics of the respective active pair: angle ϕ_1 for the right pair, and angle ϕ_3 for the left pair.

Owing to the symmetries of the SFM model, it suffices to analyze the motion of only one pair; the other is its mirror image. Figure 4(b) illustrates the kinematic analysis of the right pair. The constant d denotes the distance between the two hip joints A and B, while l is the virtual leg length. With the notation of Fig. 4(b), the (position) equation (Norton, 2008)

$$\boldsymbol{R}_{AO_1} + \boldsymbol{R}_{O_1O_2} - \boldsymbol{R}_{AB} - \boldsymbol{R}_{BO_2} = 0 \ ,$$

can be expressed as (Karydis et al, 2012)

$$le^{j(\pi-\phi_1)} + a_R e^{j(q_1-\phi_1)} - de^{j(\pi/2)} - le^{j(\phi_2)} = 0 \ , \quad (25)$$

fully determined by a single degree of freedom, the angle ϕ_1.

Each leg motion is parameterized by the leg *touchdown* and *liftoff* angles ϕ_i^{td}, and ϕ_i^{lo}, respectively (where $i = 1, \ldots, 4$), and the leg angular velocity $\dot{\phi}_{RL}$, taken to be the same for all legs. The range of motion between the touchdown and liftoff angles is captured in the *sweep* angle ψ_i, $i = 1, \ldots, 4$, and must be the same for both legs within a pair (that is, $\psi_1 = \psi_2 = \psi_R$, and $\psi_3 = \psi_4 = \psi_L$). This condition forces both legs within a pair to lift off the ground simultaneously. In order to reduce the model complexity, all touchdown angles are set at the same value ϕ^{td}. With respect to Fig. 5(a), notice that data corresponding to turning motions demonstrate a sharp change in the orientation at the beginning of each path. To capture this, we treat the initial orientation θ^{init} of the model as an additional parameter.

With these assumptions in place, the model then takes the form of a stride-to-stride map $\mathcal{M}(\xi) : \Xi \to \mathbb{R}^2 \times \mathbb{S}$ which satisfies (25) at each step, and where the parameter space is the quintuple

$$\xi = (\psi_R, \psi_L, \dot{\phi}_{RL}, \phi^{td}, \theta^{\text{init}}) \ \in \Xi \ .$$

The output of the model is taken here to be the planar position of its geometric center (x_G, y_G), that is

$$\mathsf{out}(\mathcal{M}(\xi)) = (x_G, y_G) \ .$$

B. Stochastic Extension

To capture the effect of uncertainty in the leg-ground interaction, we extend the SFM model to a stochastic setting by randomizing its parameters. In particular, the parameter vector ξ needs to become a random vector $\tilde{\xi}$, drawn according to the multivariate normal distribution

$$\tilde{\xi} \sim \mathcal{N}(\bar{\xi}, \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2, \sigma_4^2, \sigma_5^2)) \ , \quad (26)$$

where $\bar{\xi}$ is found by solving (1), and σ_i, $i = 1, \ldots, 5$ regulate the amount of noise in the model.

Selecting the model parameters over which stochasticity is introduced depends on the nature of the model, and its relation to the physical system it describes. In this particular legged example, one source of uncertainty is random leg placement, and is associated with randomizing touchdown angles (captured in σ_4). Similarly, randomization of angular velocities (captured in σ_3) is associated with speed irregularities caused by random leg-ground contact, while randomizing sweep angles (captured in σ_1, and σ_2) practically means that legs lift off the ground randomly—a phenomenon which is attributed to random leg-ground contact, and possible surface irregularities.

C. Application of the Method

We are now ready to apply the Algorithm of Table II, to estimate the model parameters for the OctoRoACH in three distinct modes of motion: (a) *straight line* (SL), (b) 90^o *clockwise turn* (CW), and (c) 90^o *counter-clockwise turn* (CCW).

With respect to Remark 2, we first select $\gamma = 0.95, \alpha = 0.29, \rho = 0.35, \varepsilon = 0.29$, and $\delta = 0.16$. Based on these selections, $N = 8$, and $M = 31$. Then, we set the length of every multisample equal to $K = 8$, which makes the total number

of experimental paths required equal to $I = M \cdot K = 248$, for each mode. Using motion capture, we collect experimental measurements of the planar position of the geometric center of the robot (x_G, y_G), and its orientation θ. The collection of K such paths generates a multisample $\mathbf{w}$. The experiments are conducted on a rubber floor mat surface, for a total time of 3 sec at a sampling rate of 20 Hz, yielding $T = 60$. The robot is set into a designated start area with an initial state set at $(x_G,\ y_G,\ \theta) = (0,\ 0,\ 0)$ [cm, cm, deg]. Initial pose errors are shown in Table III, and include measurement noise.

Figure 5(a) presents the experimental data. Let $\mathcal{W}_{\text{SL}}$, $\mathcal{W}_{\text{CW}}$, and $\mathcal{W}_{\text{CCW}}$ denote the collections of all planar trajectories of the robot for the SL, CW, and CCW control modes, respectively. The average for each set is marked $w_{\text{SL}}^{\text{ave}}$, $w_{\text{CW}}^{\text{ave}}$, and $w_{\text{CCW}}^{\text{ave}}$, respectively (shown with dashed curves). Dashed outlines represent the cone of data, $\mathsf{cone}_\gamma(\mathcal{W})$ for each case at level $\gamma = 95\%$, calculated based on $\tau_\gamma = 1.98$.[4]

Table III
INITIAL POSE ERROR STATISTICS

Type	Mean [cm cm deg]	Standard Deviation [cm cm deg]
CW	$(-0.156,\ -0.041,\ 1.23]$	$[0.177,\ 0.141,\ 1.37)$
SL	$(-0.007,\ 0.027,\ 0.06]$	$[0.234,\ 0.054,\ 1.81)$
CCW	$(-0.322,\ -0.012,\ 2.50]$	$[0.156,\ 0.130,\ 1.23)$

We first determine the deterministic component of the SFM model by solving (1) for each mode of motion, and summarize the results in Table IV. With these parameter values, the model produces paths that fit best to the experimental averages (dashed thick curves in Fig. 5(a)).

Table IV
NOMINAL SFM MODEL PARAMETERS

Type	ψ_R [deg]	ψ_L [deg]	$\dot{\phi}_{\text{RL}}$ [deg/sec]	ϕ^{td} [deg]	θ^{init} [deg]
CW	0	24.76	4.64	38.54	-16.19
SL	60.65	48.09	6.87	0.90	0
CCW	24.74	0	6.07	39.59	11.86

We investigate here the case where $\sigma_1 = \sigma_2 = \sigma_5 = 0$, assuming that only the leg placement and angular velocity are responsible for the uncertainty observed in the data—in the following Section IV-D we consider a different model parameter randomization, and compare with the results obtained in this section. Then, we follow the steps 5–11 of the Algorithm in Table II to estimate $\bar{\sigma}_3$ and $\bar{\sigma}_4$. The result of the procedure is reported in Table V.

The estimated normal distribution for the random parameter vector $\tilde{\xi}_{\text{est}}$ for each robot behavior follows from (26), with $\bar{\xi}$ shown in Table IV, $(x_G^{\text{init}}, y_G^{\text{init}}) = (0, 0)$ [cm,cm], and $\bar{\sigma}_i,\ i = 1, \ldots, 5$ found in Table V. The associated family of model instantiations $\mathcal{M}(\tilde{\xi}_{\text{est}})$ produces paths that distribute themselves over the area within the marked cone-like boundaries in Fig. 5(b) to the maximum possible degree, while allowing for a $1 - \rho = 0.65$ probability of leaving the γ-confidence region at any time t.

[4]This percentile is found according to a two-sided Student's t-distribution with 247 degrees of freedom. The 95% percentile for each multisample (with 7 degrees of freedom) is $\tau_\gamma = 2.365$.

Table V
PROBABLY APPROXIMATE NEAR MAXIMUM SFM MODEL UNCERTAINTY

Type	CW	SL	CCW
$\bar{\sigma}_1$ [deg]	0.00	0.00	0.00
$\bar{\sigma}_2$ [deg]	0.00	0.00	0.00
$\bar{\sigma}_3$ [deg/sec]	0.39	0.62	0.13
$\bar{\sigma}_4$ [deg]	17.19	4.30	18.91
$\bar{\sigma}_5$ [deg]	0.00	0.00	0.00

D. Discussion

In this section we demonstrate the applicability of the proposed framework when a different subset of model parameters is randomized. In particular, we consider the case where $\sigma_3 = \sigma_5 = 0$, assuming this time that the leg placement in both the touchdown and liftoff configurations is responsible for the uncertainty observed in the data. The parameters of the analysis remain the same as before, and we follow the steps 5–11 of the Algorithm in Table II to estimate $\bar{\sigma}_1 = \bar{\sigma}_2$ and $\bar{\sigma}_4$. The output of the procedure is reported in Table VI.

Similarly to the previous case, the random parameter vector $\tilde{\xi}_{\text{est}}$ follows from (26), with $\bar{\xi}$ shown in Table IV, $(x_G^{\text{init}}, y_G^{\text{init}}) = (0, 0)$ [cm,cm], and $\bar{\sigma}_i,\ i = 1, \ldots, 5$ tabulated in Table VI. The associated family of model instantiations $\mathcal{M}(\tilde{\xi}_{\text{est}})$ produces paths that distribute themselves over the area within the marked cone-like boundaries in Fig. 6 to the maximum possible degree, while allowing for a $1 - \rho = 0.65$ probability of leaving the γ-confidence region at any time t.

Table VI
PROBABLY APPROXIMATE NEAR MAXIMUM SFM MODEL UNCERTAINTY FOR THE SECOND PARAMETRIC RANDOMIZATION OF SECTION IV-D

Type	CW	SL	CCW
$\bar{\sigma}_1$ [deg]	8.88	4.58	8.60
$\bar{\sigma}_2$ [deg]	8.88	4.58	8.60
$\bar{\sigma}_3$ [deg/sec]	0.00	0.00	0.00
$\bar{\sigma}_4$ [deg]	0.43	4.58	0.51
$\bar{\sigma}_5$ [deg]	0.00	0.00	0.00

Comparing Fig. 6 to Fig. 5, we see that the new randomization produces paths that match the experimentally observed variability as well, albeit the predicted variability in final position for the straight line mode is more conservative. In essence, the method reports that both parameter randomization cases are acceptable solutions. Choosing one over the other ultimately relies on the designer, based on their own beliefs or assumptions on the sources of uncertainty, and which model parameters can best reflect these sources.

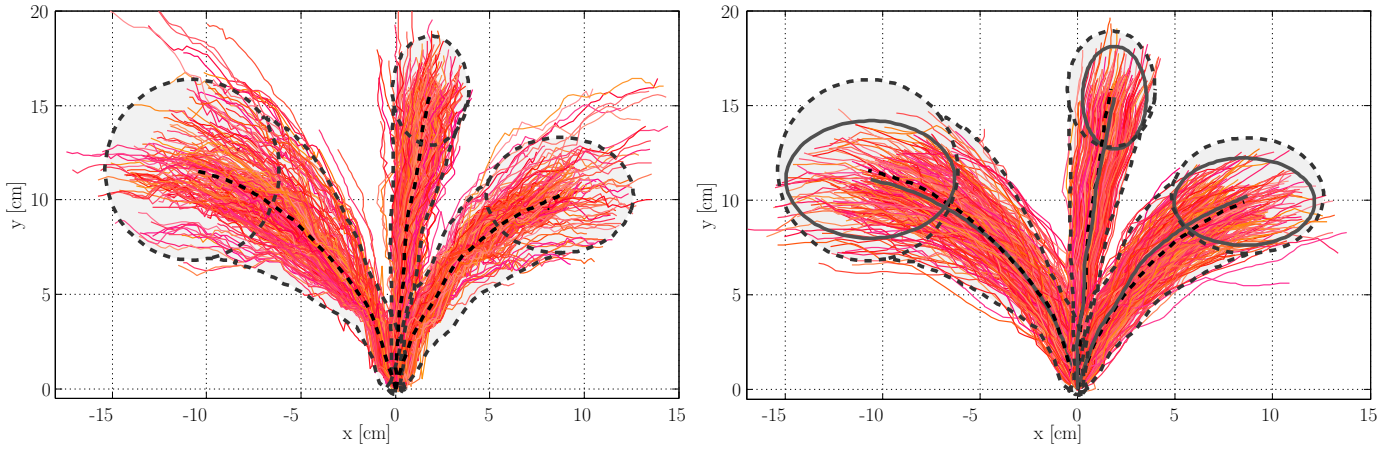

Figure 5. (a) Experimental data for the three control models considered. Dashed outlines indicate the 95% confidence interval and define the *cone of data* for each case, while experimental averages are shown with dashed thick curves in the interior of each cone. (b) Output of the stochastic model, tuned according to the values in Tables IV and V. A set of 248 random model instantiations are plotted over the experimental averages and cones of data. For all cases, the average behavior of the model, marked with a solid curve, remains very close to the experimental average (marked with a dashed curve). The 95% confidence interval at the final position (compare solid to dashed ellipses) are also matching closely.

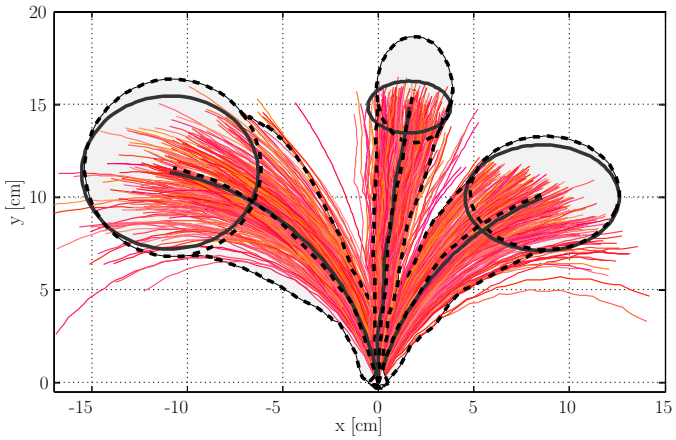

Figure 6. Output of the stochastic model, tuned according to the values in Tables IV and VI. 248 random model instantiations are plotted over the experimental averages and cones of data. Similarly to Fig. 5(b), model-predicted statistics remain very close to the experimental statistics.

V. Application to Aerial Vehicles

In the previous section we dealt with a model that is represented by a map advancing the position and orientation of the model based on a closed-form kinematic relation. The purpose of this section is to provide an example of applying the reported method to a system represented by a set of differential equations. In particular, we will use the example of a hovering task of a small-scale quadrotor and focus on its steady-state response. The performance of small-scale rotorcraft vehicles while operating in close proximity to rigid surfaces (e.g., ground, ceiling etc.) can be considerably affected by uncertain aerodynamic effects that are difficult to incorporate in low-dimensional models such as those typically used for control (Powers et al, 2012).

For clarity, here only a very small fraction of the dynamics of the physical system is excited, and thus the associated analysis allows no direct generalizations to other flying regimes; yet the case still provides an adequate example of how our method can be applied to nominal deterministic models of dynamical systems that come in the form of differential equations. Further, this example shows that the method can also work if the stochasticity is injected through an exogenous stochastic disturbance input, rather than through the model's parameters. Merely calculating a sample variance and using it directly to estimate the stochastic disturbance does not provide probabilistic guarantees of fidelity. The latter is useful for estimating what are the chances to avoid collisions. The reported method can be used to provide such guarantees.

A. The Nominal Model

In applying the method to the altitude control of small scale quadrotor, we isolate the vertical dynamics component and characterize the observed data variability at steady state. Constraining the motion along the vertical direction is achieved by the support structure of Fig. 7. A simplified quadrotor model for motion on the vertical plane and ignoring the actuator dynamics, can be expressed as (Lupashin et al, 2010)

$$\begin{aligned} m\,\ddot{y} &= -f\sin\phi \\ m\,\ddot{z} &= f\cos\phi - m\,g \\ I\,\ddot{\phi} &= \tau \end{aligned} \tag{27}$$

where f and τ are the thrust and pitch moments, respectively, m and I are the mass and moment of inertia of the vehicle, and g is the acceleration of gravity. The support structure forces $\phi = 0$, and with the addition of a gravity compensation term

$$f = m\,g + m\,u \ ,$$

the model (27) reduces to a double integrator

$$\ddot{z} = u \ . \tag{28}$$

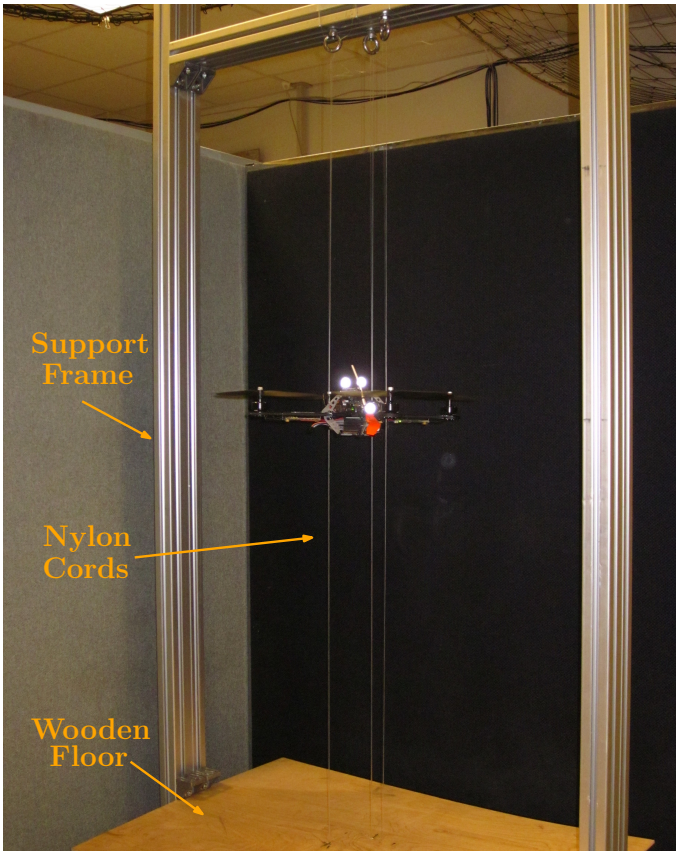

Figure 7. The experimental setup. Our support structure aids in restraining the motion along the vertical direction only. The driving strings are composed of nylon cords; they provide strong support and minimal friction, while minimizing the fluctuations on the normal to the motion plane. We also added a wooden floor to artificially generate the ground effect.

The input u is then determined by a proportional-integral-derivative (PID) controller

$$u = K_P\, e + K_I \int e\, \mathrm{dt} + K_D\, \dot{e} \ , \tag{29}$$

where $e = (r - z)$ is the position error, and r is the desired hovering height. For the purposes of this work, we consider four distinct altitudes shown in Table VII. The PID gains (K_P, K_I, K_D) are selected empirically and their values are given in Table VIII. In order to avoid unrealistic control efforts, the PID output is saturated within the region $-2 \leq u \leq 7$; the same saturation interval is used in our experiments as well.

Table VII
QUADROTOR HOVERING ALTITUDES

Case	I (Low)	II (Mid-low)	III (Mid-high)	IV (High)
r [m]	0.02	0.11	0.20	0.50

Figure 8(a) depicts the model-predicted closed-loop trajectories plotted against collected experimental data. It can be verified that the deterministic closed-loop control model is able to predict quite accurately *on average* the experimentally-observed steady-state response of the system in all four cases.

Table VIII
PID GAINS

K_P	K_I	K_D
4.00	4.50	5.00

In the following sections we will employ the Algorithm of Table II to show how the deterministic closed-loop system can be extended to a stochastic setting to capture the data variability in steady state.

B. Stochastic Extension

Contrary to the case study of Section IV, here we introduce stochasticity through an exogenous random excitation term that acts as stochastic perturbations on the nominal closed-loop dynamics (28). This shows that the method is also applicable when a system is perturbed by noise, and can be used to identify such noise terms in practice, while providing probabilistic guarantees on the reported outcome.

The closed-loop stochastic model of the system is now

$$\ddot{z} + \tilde{\xi} = u \ ,$$

where u is calculated by the same PID controller as in (29) and then saturated in the region $-2 \leq u \leq 7$, with $\tilde{\xi} \sim \mathcal{N}(0, \sigma_\xi^2)$, and the variance σ_ξ^2 to be determined by the Algorithm of Table II. In essence, $\tilde{\xi}$ here corresponds to a zero-mean Gaussian process corrupting the control effort, and the task of the method is to find the variance σ_ξ^2 that captures the data variability in steady state under prespecified fidelity specifications. Figure 9 summarizes schematically the closed-loop control stochastic model we consider here.

C. Application of the Method

In this section we apply the proposed methodology to determine σ_ξ^2 at *steady state* for each of the four cases shown in Table VII. We select the same problem parameters as in the previous case study: $\gamma = 0.95, \alpha = 0.29, \rho = 0.35, \varepsilon = 0.29$, and $\delta = 0.16$, giving $N = 8$, $M = 31$ with $K = 8$. We collect closed-loop altitude measurement data from a total of 248 experimental trajectories, for each height. Trials for the first three cases last 15 sec; for the fourth, 20 sec, because the system enters its steady state at the end of the 15^{th} second. The feedback loop refreshes at 30 Hz, at the same frequency as our motion capture system. The initial pose errors are not significant since the support structure ensures that the start position remains unchanged in all trials.

Figure 8(a) presents the experimentally observed trajectories. Let $\mathcal{W}_{\text{l}}$, $\mathcal{W}_{\text{ml}}$, $\mathcal{W}_{\text{mh}}$, $\mathcal{W}_{\text{h}}$ indicate the collections of the trajectories for the low, mid-low, mid-high, and high hovering altitude, respectively. Let w denote an element in these sets. The average for each case is denoted by $w_{\text{l}}^{\text{ave}}$, $w_{\text{ml}}^{\text{ave}}$, $w_{\text{mh}}^{\text{ave}}$, $w_{\text{h}}^{\text{ave}}$, respectively, shown in solid thick curves. Dashed curves mark the cone of data for each case, which is calculated based on a two-sided Student's t-distribution with 247 degrees of freedom, the 95th percentile of which is $\tau_\gamma = 1.98$. The

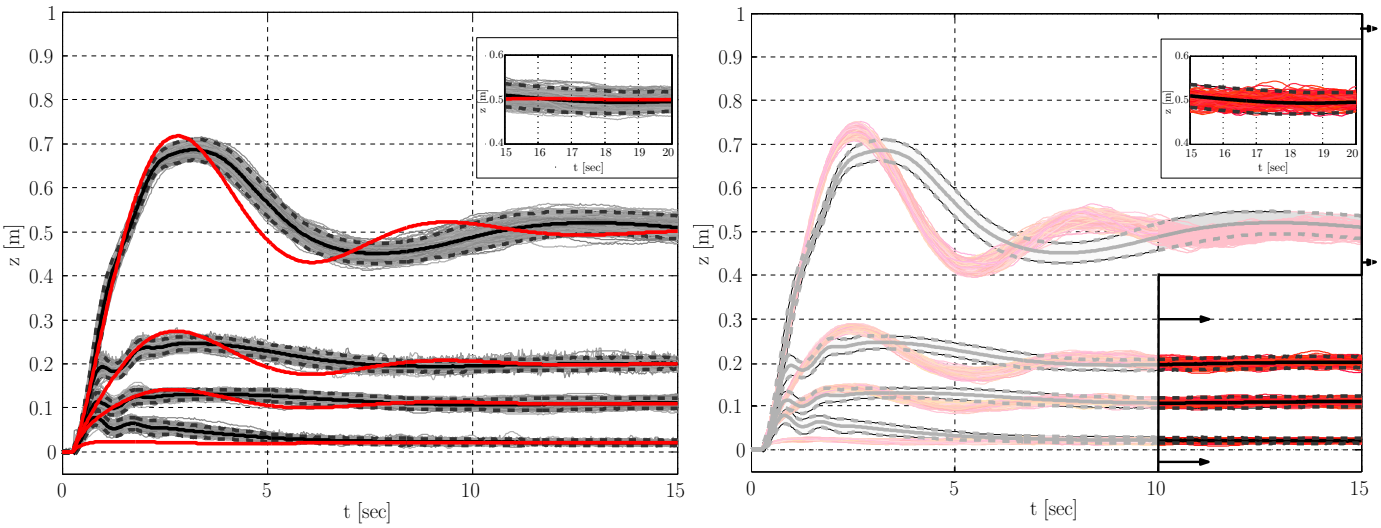

Figure 8. (a) Experimentally collected paths for a quadrotor aerial vehicle tasked to hover at four distinct altitudes (0.02 m, 0.11 m, 0.20 m, and 0.50 m, respectively). Individual paths are enclosed within the respective cone-of-data areas (marked with the dashed curves), while the experimental averages are shown in solid thick curves. For the fourth case, the steady state is achieved at $t = 20$ sec. For clarity purposes, we show the last 5 sec in the add-on window on the top right corner of the figure. The superimposed solid thick curves correspond to the model-predicted outputs according to (29). (b) The output of the stochastically perturbed control architecture depicted in Fig. 9, where the values of σ_ξ for each case have been estimated by the proposed framework, and are shown in Table IX. A set of 248 random model instantiations are plotted on top of the experimental averages and cones of data. We are now able to capture the data variability during the *steady-state* response.

additional 5 sec for the highest altitude case are shown as an add-on figure on top of Fig. 8(a).

The results of the application of the proposed method are summarized in Table IX. Figure 8(b) depicts 248 randomly generated model paths for each case, parameterized according to the values in Tables VIII and IX. Focusing on the steady-state response of the system only, we can verify that the stochastically extended control scheme of Fig. 9 is capable of capturing the data variability, and that the resulting paths match closely their experimental counterparts. We highlight in Fig. 8(b) the steady-state part of system responses to emphasize that our focus is in this particular regime. The induced ground effect is evident in the behavior of the quadrotor closest to the ground, where it appears as if the generated airflow creates an aerodynamic "cushion" below the platform. This has a stabilizing effect on the platform, as indicated by the reduction in the amplitude of the vehicle's residual oscillations at steady state.

Table IX
PROBABLY APPROXIMATE NEAR MAXIMUM QUADROTOR MODEL UNCERTAINTY

Case	I (Low)	II (Mid-low)	III (Mid-high)	IV (High)
$\bar{\sigma}_\xi$	0.08	0.13	0.17	0.30

Note that the control architecture of Fig 9 does not capture the ripples during the transient phases. These ripples could be attributed, at least in part, to the nonlinear coupling between the pitch angle ϕ and the actuator input f (cf. (27)) which are currently ignored by assuming that $\phi \equiv 0$. Thus the simplified deterministic model is incapable of reproducing this transient,

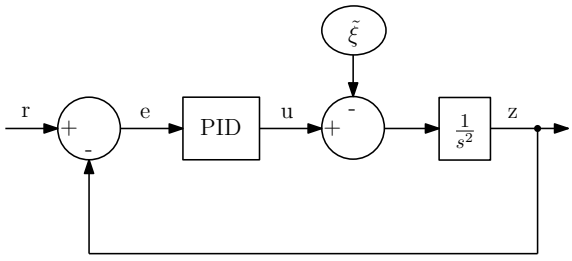

Figure 9. Schematic representation of the closed-loop control stochastic model for the steady-state vertical dynamics component during quadrotor hovering. The input r denotes the desired hovering height, while the gains of the PID controller have been tuned a-priori, and remain the same with those used in our data collections. In $\tilde{\xi} \sim \mathcal{N}(0, \sigma_\xi^2)$, the variance σ_ξ^2 is estimated using the proposed algorithm.

irrespectively of how stochasticity is infused. Using a full nonlinear model such as (27) may enable one to capture this transient behavior, however doing so falls outside the scope of the present paper.

D. Discussion

This section discusses an alternative way to perturb the nominal closed-loop dynamics (28), and provides some insight on the predictive ability of the augmented stochastic architecture shown in Fig. 9.

We first perturb the output of the model with an exogenous random term—modeled again as a zero-mean Gaussian process process—shown in Fig. 10.[5] Then, we apply the proposed method with all parameters retaining the same values

[5]Infusing stochasticity in this way may also have practical significance in the case of uncertain or noisy state measurements.

as in Section V-C. The results are summarized in Table X, and Fig. 11 depicts 248 randomly generated model paths for each case, parameterized according to the values found in Tables VIII, and X. The modified stochastic extension also captures the data variability once the steady state has been reached, but system responses are somewhat different (qualitatively) from the experimental data due to "chattering" (see Fig. 11). Overall, however, the methodology does not promote any one option over another; instead, it reports that both are possible solutions. The designer has then to select which solution to keep, based on their own beliefs or assumptions on the possible sources of uncertainty.

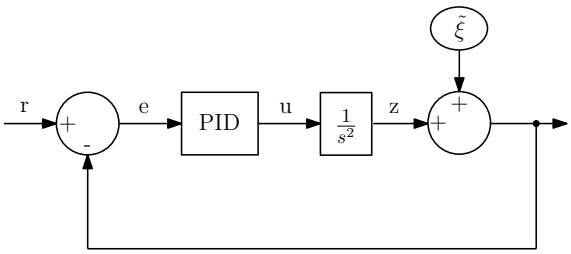

Figure 10. Schematic representation of the modified closed-loop control architecture. As before, the input r denotes the desired hovering height, and the PID gains have been tuned a-priori.

Table X
PROBABLY APPROXIMATE NEAR MAXIMUM QUADROTOR MODEL UNCERTAINTY WHEN THE EXCITATION TERM AFFECTS THE MODEL OUTPUT

Case	I (Low)	II (Mid-low)	III (Mid-high)	IV (High)
$\bar{\sigma}_\xi$	0.002	0.004	0.005	0.007

Next, we examine how the stochastic architecture of Fig. 9, with variance parameter σ_ξ^2 specified by data corresponding to a single desired hovering altitude, can be used to predict the steady-state response of the vehicle both temporally, and across different operating points. In detail, we use as training data those corresponding to case II (i.e., $z_{\text{train}} = 0.11$ m), with $\bar{\sigma}_{\xi,\text{train}} = 0.13$ (from Table IX). Then, we use the same model, with the same statistics for randomized exogenous forces, but with different inputs (desired hovering altitudes) to predict the steady-state responses in the remaining cases of Table VII.

Figure 12 presents the results. As expected, the data variability observed in steady state is captured faithfully when the desired reference point does not change too much (i.e., $z_{\text{pred}} = 0.2$ m). However, when the desired reference point induces different operating conditions in which the nature of the environmental interaction with the physical system changes, the resulting stochastic model may offer a rough prediction, albeit collecting more experimental data for these set points is recommended. For example, when the set point is $z_{\text{pred}} = 0.02$ m, we get a more conservative picture of what is happening since the ground effect becomes evident, manifesting itself as a form of a stabilizing aerodynamic cushion under the vehicle (compare with Fig. 8).

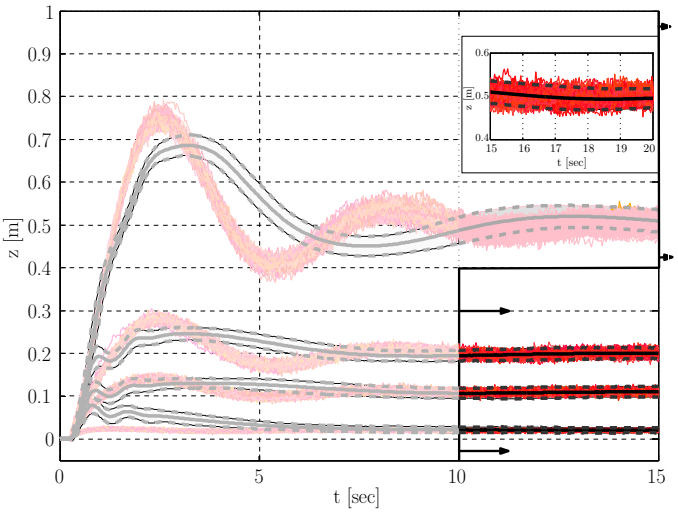

Figure 11. The output of the stochastically perturbed control architecture depicted in Fig. 9, where the values of σ_ξ for each case have been estimated by the proposed framework, and are shown in Table IX. A set of 248 random model instantiations are plotted against the experimental averages and cones of data.

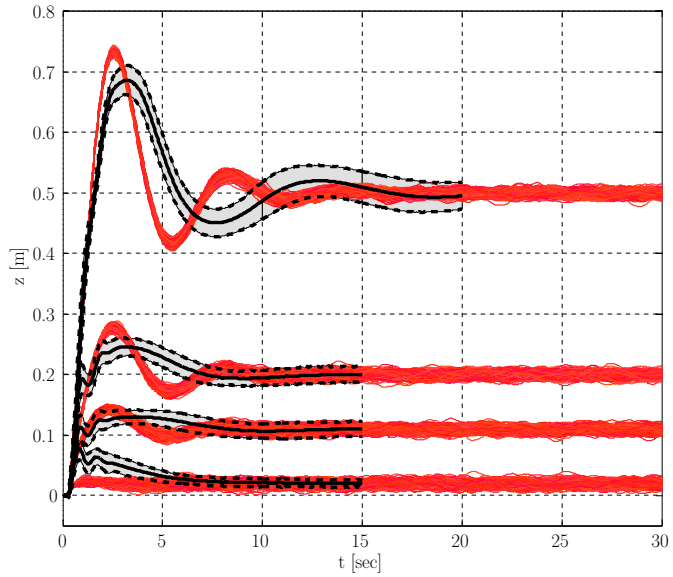

Figure 12. Training data and exogenous excitation statistics correspond to the second case only ($z_{train} = 0.11$ m, and $\bar{\sigma}_{\xi,train} = 0.13$). The stochastically extended control architecture (Fig. 9) is then used to predict the steady-state response of the system under different inputs, as well as extend these predictions temporally. Provided that the operating conditions do not vary significantly, the stochastic extension is able to make accurate predictions, both temporal, and for different reference points.

The same setup is also used to make temporal predictions that extend for additional 15 s for cases I – III, and 10 s for case IV. As before, provided that the induced operating conditions do not vary significantly, the stochastically extended architecture is able to make accurate temporal predictions, with the caveat that these predictions come with no probabilistic guarantees.

VI. Conclusions and Discussion

The paper reports on a data-driven framework for extending existing deterministic models into a stochastic regime, while providing probabilistic guarantees on the reported outcome. This can be useful in cases where an otherwise adequate model of a system is available, but may fail to capture possible uncertain system-environment interactions. The reported approach combines methodological elements of probabilistic model validation and randomized algorithms to simultaneously quantify the fidelity of a model, and tune the distribution of random parameters in the augmented stochastic extension to enable it to reproduce the variability observed experimentally in a physical process of interest. The approach is independent of the type and structure of the model considered, and can be applied to a variety of deterministic and stochastic models, including ones based on differential equations.

We demonstrate the applicability of the method in two examples involving very different base models. The first system expresses the stochastic kinematics of the horizontal-plane motion of the miniature legged robot OctoRoach in the form of a single-parameter stride-to-stride map; the second system is a second order stochastically perturbed differential equation modeling the vertical flight of a quadrotor close to the ground during its steady state. The applicability of the approach to such a wide range of system representations offers promise in terms of being potentially useful across domains.

Due to the existence of the underlying deterministic model on which the stochastic extension is based, the latter can be used to make predictions under different circumstances for which no training data are available. However, the predictive ability of the stochastic model should be judged within the resolution constraints of the underlying deterministic model. As Fig. 12 indicates, the stochastically extended model will succeed in predicting new system behaviors accurately when new circumstances do not alter significantly the nature of the operating conditions, at the expense that the provided probabilistic guarantees will no longer hold. In a different regime of operating conditions, part of the predictive ability of the (combined) stochastic extension can be carried over to the new parameter regime, with the caveat, however, that accurate characterization really requires additional experimental data.

There are multiple ways to infuse stochasticity into the underlying deterministic model. In this paper, for example, we have chosen to capture the leg-ground uncertain interaction of the legged vehicle case by infusing stochasticity in its model parameters, while in the aerial vehicle case, uncertain aerodynamic effects corrupt the control input signal, and noisy measurements can be captured by randomizing the output signal. The methodology described does not promote any one option over another, but rather offers the mechanism for generating the different options. Ultimately, the designer needs to choose the way to infuse stochasticity, and decide which resulting model output is preferable; their insight into the possible sources of uncertainty affecting the system and the mechanisms through which the uncertainty affects the system at hand, provides the basis for such choices.

Funding

This work was supported by National Science Foundation (grant numbers IIS-1350721, CMMI-1130372, and CNS-1035577), and the Army Research Laboratory (grant number W911NF-08-2-0004).

References

Abbeel P, Coates A, Montemerlo M, Ng AY, Thrun S (2005) Discriminative training of kalman filters. In: Thrun S, Sukhatme GS, Schaal S (eds) Robotics: Science and Systems I, MIT Press, pp 289–296

Abu-Mostafa YS, Magdon-Ismail M, Lin HT (2012) Learning From Data. AMLBook

Alamo T, Tempo R, Camacho E (2009) Randomized Strategies for Probabilistic Solutions of Uncertain Feasibility and Optimization Problems. IEEE Transactions on Automatic Control 54(11):2545–2559

Alterovitz R, Branicky M, Goldberg K (2008) Motion Planning Under Uncertainty for Image-guided Medical Needle Steering. The International Journal of Robotics Research 27(11-12):1361–1374

Anderson R, Milutinović D (2014) A Stochastic Optimal Enhancement of Feedback Control for Unicycle Formations. Robotica 32(2):305–324

Aoude G, Luders B, Joseph JM, Roy N, How JP (2013) Probabilistically Safe Motion Planning to Avoid Dynamic Obstacles with Uncertain Motion Patterns. Autonomous Robots 35:51–76

van den Berg J, Abbeel P, Goldberg K (2011) LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. The International Journal of Robotics Research 30(7):895–913

Blackmore L, Ono M, Williams B (2011) Chance-Constrained Optimal Path Planning with Obstacles. IEEE Transactions on Robotics 27(6):1080–1094

Calafiore G, Dabbene F, Tempo R (2011) Research on Probabilistic Methods for Control System Design. Automatica 47(7):1279–1293

Censi A, Calisi D, De Luca A, Oriolo G (2008) A bayesian framework for optimal motion planning with uncertainty. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 1798–1805

Chirikjian GS (2012) Stochastic Models, Information Theory, and Lie Groups, Volume 2: Analytic Methods and Modern Applications. Birkhauser

Cizelj I, Belta C (2014) Control of noisy differential-drive vehicles from time-bounded temporal logic specifications. The International Journal of Robotics Research 33(8):1112–1129

Dabbene F, Sznaier M, Tempo R (2012a) A probabilistic approach to optimal estimation - part II: Problem formulation and methodology. In: Proceedings of the 51th IEEE Conference on Decision and Control, pp 196–201

Dabbene F, Sznaier M, Tempo R (2012b) A probabilistic approach to optimal estimation part I: Problem formulation and methodology. In: Proceedings of the 51th IEEE Conference on Decision and Control, pp 190–195

Deisenroth MP, Rasmussen CE (2011) PILCO: A model-based and data-efficient approach to policy search. In: Proceedings of the International Conference on Machine Learning

Dixon WJ, Massey FJ (1957) Introduction to Statistical Analysis. McGraw Hill

Dubins L (1957) On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. American Journal of Mathematics 79(3):497 – 516

Efron B, Tibshirani R (1994) An Introduction to the Bootstrap. CRC Press

Gevers M, Bombois X, Codrons B, Scorletti G, Anderson BD (2003) Model Validation for Control and Controller Validation in a Prediction Error Identification Framework – Part I: Theory. Automatica 39(3):403–415

Gillula JH, Hoffmann GM, Huang H, Vitus MP, Tomlin CJ (2011) Applications of hybrid reachability analysis to robotic aerial vehicles. The International Journal of Robotics Research 30(3):335–354

Halder A, Bhattacharya R (2014) Probabilistic Model Validation for Uncertain Nonlinear Systems. Automatica 50(8):2038–2050

Hall J, Rasmussen CE, Maciejowski J (2012) Modelling and control of nonlinear systems using gaussian processes with partial model information. In: Proceedings of the 51st IEEE Conference on Decision and Control, pp 5266–5271

Haykin S (1999) Neural Networks: A Comprehensive Foundation. Prentice-Hall, Upper Saddle River, NJ

Hofmann T, Schölkopf B, Smola AJ (2008) Kernel Methods in Machine Learning. Annals of Statistics 36(3):1171–1220

Holmes P, Full RJ, Koditschek DE, Guckenheimer J (2006) The Dynamics of Legged Locomotion: Models, Analyses, and Challenges. SIAM Review 48(2):207–304

Hoover AM, Burden S, Fu XY, Sastry S, Fearing RS (2010) Bio-inspired design and dynamic maneuverability of a minimally actuated six-legged robot. In: Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics, pp 869–876

Kalakrishnan M, Chitta S, Theodorou E, Pastor P, Schaal S (2011) STOMP: Stochastic trajectory optimization for motion planning. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 4569–4574

Karydis K, Poulakakis I, Tanner HG (2012) A switching kinematic model for an octapedal robot. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp 507–512

Karydis K, Poulakakis I, Tanner HG (2013) Probabilistic validation of a stochastic kinematic model for an octapedal robot. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 2562–2567

Karydis K, Zarrouk D, Poulakakis I, Fearing R, Tanner H (2014) Planning with the STAR(s). In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp 3033–3038

Karydis K, Liu Y, Poulakakis I, Tanner HG (2015) A template candidate for miniature legged robots in quasi-static motion. Autonomous Robots 38(2):193–209

Kewani G, Ishigami G, Iagnamma K (2009) Stochastic mobility-based path planning in uncertain environments. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp 1183–1189

Koltchinskii V, Abdallah C, Ariola M, Dorato P, Panchenko D (2000) Statistical learning control of uncertain systems: it is better than it seems. Tech. Rep. EECE-TR-00-001, University of New Mexico

Kumar V, Michael N (2012) Opportunities and challenges with autonomous micro aerial vehicles. The International Journal of Robotics Research 31(11):1279–1291

Lee L, Poolla K (1996) On Statistical Model Validation. Journal of Dynamic Systems, Measurement, and Control 118(2):226–236

Ljung L (1999) System Identification: Theory for the User. Prentice-Hall

Long AW, Wolfe KC, Mashner M, Chirikjian GS (2012) The Banana Distribution is Gaussian: A Localization Study with Exponential Coordinates. In: Proceedings of Robotics: Science and Systems

Lupashin S, Schollig A, Sherback M, D'Andrea R (2010) A simple learning strategy for high-speed quadrocopter multi-flips. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 1642–1648

Mellinger D, Michael N, Kumar V (2012) Trajectory generation and control for precise aggressive maneuvers with quadrotors. The International Journal of Robotics Research 31(5):664–674

Mendelson S (2003) A Few Notes on Statistical Learning Theory. In: Mendelson S, Smola EJ (eds) Advanced Lectures in Machine Learning, vol 2600, Springer Verlag, pp 1–40

Murphy KP (2012) Machine Learning: A Probabilistic Perspective. The MIT Press

Norton R (2008) Design of Machinery. McGraw Hill, New York, NY

Ogunfunmi T (2007) Adaptive Nonlinear System Identification: The Volterra and Wiener Model Approaches. Signal and Communication Technology, Springer-Verlag, New York, NY

Pereira AA, Binney J, Hollinger GA, Sukhatme GS (2013) Risk-aware Path Planning for Autonomous Underwater Vehicles using Predictive Ocean Models. Journal of Field Robotics 30(5):741–762

Pivtoraiko M, Mellinger D, Kumar V (2013) Incremental micro-UAV motion replanning for exploring unknown environments. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 2452–2458

Powers C, Mellinger D, Kushleyev A, Kothmann B, Kumar V (2012) Influence of aerodynamics and proximity effects in quadrotor flight. In: Proceedings of the International Symposium on Experimental Robotics, Springer Tracts in Advanced Robotics, vol 88, pp 289–302

Prajna S (2006) Barrier Certificates for Nonlinear Model Validation. Automatica 42(1):117–126

Proctor J, Holmes P (2008) Steering by Transient Destabilization in Piecewise-Holonomic Models of Legged Locomotion. Regular and Chaotic Dynamics 13(4):267–282

Pullin A, Kohut N, Zarrouk D, Fearing R (2012) Dynamic turning of 13 cm robot comparing tail and differential drive.

In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 5086–5093

Qian F, Zhang T, Li C, Masarati P, Hoover A, Birkmeyer P, Pullin A, Fearing R, Goldman D (2012) Walking and runnign on yielding and fluidizing ground. In: Proceedings of Robotics: Science and Systems, pp 345–352

Rasmussen CE, Williams CKI (2006) Gaussian Processes for Machine Learning. MIT Press

Ray L, Stengel R (1993) A Monte Carlo Approach to the Analysis of Control System Robustness. Automatica 29(1):229–236

Schetzen M (2006) The Volterra and Wiener Theories of Nonlinear Systems. Krieger Publishing, Malabar, FL

Schmidt M, Lipson H (2009) Distilling Free-Form Natural Laws from Experimental Data. Science 324:81–85

Seegmiller N, Rogers-Marcovitz F, Miller G, Kelly A (2013) Vehicle model identification by integrated prediction error minimization. The International Journal of Robotics Research 32(8):912–931

Shah S, Pahlajani C, Lacock N, Tanner H (2012) Stochastic receding horizon control for robots with probabilistic state constraints. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 2893–2898

Shkolnik A, Levashov M, Manchester IR, Tedrake R (2011) Bounding on rough terrain with the LittleDog robot. The International Journal of Robotics Research 30(2):192–215

Steinhardt J, Tedrake R (2012) Finite-time regional verification of stochastic non-linear systems. The International Journal of Robotics Research 31(7):901–923

Tempo R, Calafiore G, Dabbene F (2012) Randomized Algorithms for Analysis and Control of Uncertain Systems: With Applications, 2nd edn. Springer Publishing Company, Inc.

Thrun S, Burgard W, Fox D (2005) Probabilistic Robotics (Intelligent Robotics and Autonomous Agents). The MIT Press

Timcenko A, Allen P (1993) Modeling dynamic uncertainty in robot motions. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 531–536

Vapnik VN (1998) Statistical Learning Theory. Wiley, New York, NY

Vidyasagar M (2001) Randomized Algorithms for Robust Controller Synthesis using Statistical Learning Theory. Automatica 37(10):1515–1528

Vidyasagar M (2003) Learning and Generalization With Applications to Neural Networks, 2nd edn. Springer-Verlag, London, UK

Zarovy S, Costello M, Mehta A, Gremillion G, Miller D, Ranganathan B, Humbert JS, Samuel P (2010) Experimental study of gust effects on micro air vehicles. In: AIAA Conference on Atmospheric Flight Mechanics., pp AIAA–2010–7818

Zarrouk D, Pullin A, Kohut N, Fearing R (2013) STAR - a sprawl tuned autonomous robot. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp 20–25