

Data-driven Abstractions for Robots with Stochastic Dynamics

Herbert G. Tanner and Adam Stager

Abstract—This paper describes the construction of stochastic, data-based discrete abstractions for uncertain random processes continuous in time and space. Motivated by the fact that modeling processes often introduce errors which interfere with the implementation of control strategies, here the abstraction process proceeds in reverse: the methodology does not abstract models; rather it models abstractions. Specifically, it first formalizes a template for a family of stochastic abstractions, and then fits the parameters of that template to match the dynamics of the underlying process and ground the abstraction. The paper also shows how the parameter fitting approach can be implemented based on a probabilistic model validation approach which draws from randomized algorithms, and results in a discrete abstract model which is approximately simulated by the actual process physics, at a desired confidence level. In this way, the models afford the implementation of symbolic control plans with probabilistic guarantees at a desired level of fidelity.

Index Terms—stochastic processes; discrete abstractions; simulation relations; randomized algorithms

I. INTRODUCTION

There are multiple robotic systems the evolution of which exhibits *stochasticity* and levels of noise that cannot be ignored. Examples of such systems are micro aerial vehicles (especially when flying near surfaces and are affected by ground effects) and miniature legged robots [1], [38], ocean drifters [2], and micro robots navigating in solution [3]. The motion of robotic systems such as these is better modeled using continuous or discrete stochastic processes. Stochastic processes offer an important modeling framework for dynamic phenomena that exhibit uncertainty [4]. Indeed, the flexibility and expressivity of stochastic models affords them applications in a diverse range of fields, from engineering [4], to physics and biology [5], to finance [6]. Yet, stochastic processes are notoriously difficult to analyze and design control laws for [7].

Uncertain physical processes are encountered in the study of complex phenomena such as legged locomotion on granular media [8], and ground-effect aerodynamic coupling in quadrotor flight [9]. Detailed modeling of such interactions not only requires deep understanding of the underlying physics but leads to mathematical descriptions which may be too cumbersome for planning and control design purposes. This motivates the construction and use of simplified abstract representations of the effect of these complex physical processes on the motion behavior of systems. Some of these abstract representations capture uncertainty in a deterministic way; for example in the form of some slowly-varying disturbance [9], or by having a

range of case-specific models [10]. Other times, uncertainty is being modeled in the form of stochastic noise—a formulation that yields compact mathematical representations [11] in the form of stochastic differential equation (SDE)s. These continuous in time and space mathematical models may be elegant and concise, but they are notoriously difficult to analyze in closed form in multiple dimensions [5]. A question that arises is whether these models admit further simplification that makes them amenable to analysis and design, without depriving them of the ability to make reliable predictions of future behavior.

In this light, symbolic control synthesis methods with their automated tools for property verification and control synthesis, appear particularly attractive. There is a mature body of literature offering a variety of algorithms, primarily based on fixed-point computations [12], that operate on discrete abstractions of dynamical systems to yield controllers that enforce linear temporal logic (LTL) specifications [13], [14]. If a discrete abstraction is linked to its concrete system formally through a simulation or a bisimulation relation (i.e. [15]–[17]; for exceptions, see [18]), then model checkers can be used for the verification of LTL specifications, and symbolic controllers can be refined to regulate the concrete system. All these formal-methods tools for verification and synthesis rely on the availability of discrete models for the processes at hand.

The application of formal methods to the analysis of *stochastic* processes follows a parallel track [19]–[22], typically leveraging tools for non-probabilistic systems [23]; however, existing literature on abstracting stochastic processes has a common starting point: a *known* model of the concrete stochastic process, usually in the form of an SDE or a Markov decision process (MDP) [21], [22]. For instance, there can be a finite-state probabilistic system [24], or an MDP [23], where the transition probabilities are given, or an SDE with known diffusion and drift coefficients [25], [26]. In this context, emphasis has been given on methods that alleviate the curse of dimensionality by adapting [21] or avoiding [26] state and input discretization, usually by exploiting stability properties—although forward completeness [27] *does* draw from input-to-state stability concepts. Still, the abstraction process starts with a *fully known* concrete model for the stochastic dynamics.

This paper reports on an approach that uses *experimental data* from the stochastic phenomenon of interest to construct a discrete abstraction of the process. There is a pathway that formally links the underlying physics to abstract models—phenomenological or based on first-principles—in a way that is useful for planning, control, and verification. The approach bears some similarities to existing work [21], [22], in the sense of producing a Markov chain as a discrete abstraction of the concrete stochastic process. In order, however, for these low-

Bert Tanner and Adam Stager are with the Department of Mechanical Engineering, University of Delaware. Email: {btanner,astager}@udel.edu

Special thanks to Brad Hobbs for his assistance in data analysis, and to Andreas Malikopoulos for sharing his scaled city environment.

dimensional representations to make predictions that match experimental observations, the former have to be populated with parameter values based on *data*. Thus, among the key conceptual differences of this paper to the aforementioned work is that it works with data instead of a concrete (hybrid system) model for the underlying physical process, and does not place the emphasis on the behaviors of the models being “close,” but rather to one system’s behavior being a subset of that of the other.

An application area where the ability to predict system behavior based on a simple stochastic model is particularly useful, is planning the motion of miniature legged robots [28], [29]. First-principle models for multi-legged robots are difficult to obtain, complex, and dependent on a variety of partially known parameters. The lack of reliable simple models for this class of robots impedes efforts to analyze their motion behavior and design motion controllers. The significance of having a model which cannot exhibit behavior foreign to the stochastic process it represents, is that one can confidently predict and reason about future behavior and performance, under different operating conditions, without additional explicit experimentation over the whole operating regime.

The paper marks a departure from the norm in the context of deriving finite abstractions, for cases where data are available about the behavior of the physical process that is being studied. The reason for this departure is that existing techniques inadvertently propagate modeling errors through the abstraction process and eventually leave them unchecked. The hypothesis here is that the fidelity of the abstraction can be ensured during modeling process, and uncertainty from both process and modeling be probabilistically quantified. So contrary to the prevailing trend in existing literature which *abstracts models* of physical processes, this approach advocates *modeling an abstraction* of the uncertain process. For example, if experimental data are available, the typical approach is to first fit a continuous concrete model to these data, and then derive a finite abstraction for that model; alternatively, what is done here is that some uncertain finite abstraction is formalized first, and then this abstraction is related to the concrete finite model based on the data. Note that related literature reserves the terms “concrete” and “abstract” for detailed high-dimensional, and simplified low-dimensional models, respectively; here, however, the “abstraction” is low-dimensional but not “concrete” because it is uncertain, while the “model” is of the same dimensionality as the abstraction, but it is concrete in the sense of being completely known.

The paper’s contribution is in (i) determining conditions under which a discrete model of a stochastic process can be related through a simulation relation to the stochastic process; (ii) introducing the concept of (approximate) strong simulation relations between discrete stochastic processes and a means of practically establishing them through a data-driven approach, and (iii) demonstrating the use of the relations on predicting the behavior of uncertain miniature legged robots in planar motion planning tasks. A corollary of this analysis is that in order for a model to *partially* capture the behavior of the underlying physics, it has to exhibit more uncertain behavior compared to the physics it attempts to model.

The methodology combines simulation relations with randomized algorithms to produce data-driven probabilistic discrete abstractions for stochastic processes lacking a concrete model. It does so in a way that the abstract model can enjoy formal fidelity guarantees with respect to the physical processes it describes. Given the simulation relation established between the *discrete* probabilistic abstractions and its underlying physical process, the resulting model can feed into existing tools for verification and synthesis [19]–[22]. Using a body of experimental data, the performance certificates established for the abstract model can now be extended to the physics being modeled—at a lower level than that of the SDE or the MDP typically used as a concrete model.

II. PROBLEM STATEMENT

Terminology and notation follows Rogers and Williams [30]. Existing concepts and established definitions from stochastic processes and formal languages are provided in the Appendix, and whenever the concepts are invoked, a reference to the corresponding definition in the Appendix is provided in a footnote. Whatever stochastic processes-related symbols and concepts are not defined explicitly in this paper, can be found in Rogers and Williams [30]. A table of the notation symbols used in this paper is provided before the Appendix.

The objective is to construct formal abstractions of stochastic physical processes, in the form of discrete probabilistic models. To establish a starting point for this process, it is first assumed that the physical phenomenon of interest is adequately captured by the mathematical model of a (continuous-space) Markov process¹ evolving on some subset E of a Euclidean space.

The phenomenon of interest is thus captured by a Markov process X . The process X is in fact a special mapping from Ω to a function space E^T : the space of functions which map into E , and depend on a variable which takes values in T . One then writes $X : \Omega \rightarrow E^T$, $\omega \mapsto X(\omega)$. (One example of T is a time domain, a subset of $\mathbb{R}_{\geq 0}$.) If the *valuation map* on functions, $\pi_t : E^T \rightarrow E$, is defined as $\pi_t(f) := f(t)$, and the σ -algebra² \mathcal{E}^T on E is expressed as $\mathcal{E}^T := \sigma\{\pi_t : t \in T\}$, then the function $X(\omega)$ is understood as an (E^T, \mathcal{E}^T) -random variable, also known as a sample path.

Imagine that X is parameterized by control signals, which are assumed to be finite functions of time $t \in T$, and of state $x \in E$. This setup gives rise to a *controlled* continuous Markov process (not be confused with a Markov decision process, which is discrete in space and time and also involves rewards and discount factors). For the particular class of controlled Markov decision processes considered here, the space of control laws is finite. Essentially, the controlled Markov processes of this paper are finite families of (closed-loop) continuous Markov processes.

With a controlled continuous-space Markov process capturing the dynamics of the phenomenon of interest, the question now is two-fold: (i) what type of discrete, in space and time,

¹See Definition 13 in Appendix B.

²See Definition 8 in Appendix B.

models can serve as abstractions of such stochastic processes, and (ii) under what conditions can one relate formally such abstractions to the original processes, in a way that guarantees that *the behavior of the abstraction is necessarily exhibited by the process* too. The shows that relaxed versions of the formal relations between abstraction and process can be still be established with the same performance properties but at a reduced (probabilistic) level of confidence, if due to lack of information or derivational and computational issues, doing so analytically becomes impractical.

III. APPROACH

If a discrete probabilistic model of some stochastic physical process is constructed in a way that *conservatively* represents the original physical process, then a strong simulation relation can be established between process and model, ensuring that model behavior is a subset of that of the stochastic process. This inclusion property ensures that control strategies devised based on the model can be refined on the physical system. The discussion starts with the assumption that the controlled continuous-space Markov processes is an accurate and full representation of the actual physical phenomenon.

In order to make these controlled continuous Markov processes amenable to analysis using tools from formal languages, this paper applies one form of discretization. Both the space of control signals, and the Markov process's state space is discretized. Control discretization is more straightforward, so the discussion begins with that.

Recall that the space of control laws driving X_t has been assumed to be a finite collection of bounded functions on $\mathbb{R}_+ \times E$. To every such function, associate a distinct symbol σ and collect all the symbols in a finite set Σ .

The discretization of the state space of X_t is a little more involved. It is induced by imposing a certain structure on \mathcal{E} . Imagine “seeding” E with N isolated points, and define *disjoint* compact neighborhoods $\Gamma_1, \dots, \Gamma_N$ around those points. Let $\mathbb{1}_{\Gamma_i}$ denote the indicator function of neighborhood Γ_i . Define $\Gamma_{N+1} := E \setminus \bigcup_{i=1}^N \Gamma_i$, and (re)define \mathcal{E} to be the σ -algebra generated³ by the sets Γ_i in the form $\mathcal{E} := \sigma\{\Gamma_1, \dots, \Gamma_N, \Gamma_{N+1}\}$. (Mind the subtle appearance distinction between σ , the control symbol, and σ , the σ -algebra.)

Now fix $\sigma \in \Sigma$, and consider sample paths X_t as collections of (E, \mathcal{E}) -random variables (or equivalently, X as a (E^T, \mathcal{E}^T) -random variable), where specifically $T = [0, \tau] \subset \mathbb{R}_+$ for fixed $\tau > 0$. If, for example, one considers the path of a robot within its workspace produced with some fixed control law σ , these sample paths will correspond to random realizations of the robot's evolution over a time period $[0, \tau]$. For $n \in \mathbb{N}$, the finite collection of (E^T, \mathcal{E}^T) -random variables $X_{n\tau}$ now defines a (discrete-time) Markov process $\{X_n : n \in \mathbb{N}\}$. Continuing with the aforementioned analogy, this Markov process can be viewed as a probabilistic reachability graph—with the caveat that the robot's control law is fixed. By virtue of Γ_i being disjoint and each X_n being a function, at time t only a single

event (e.g., “the robot is at workspace cell Y ”) can occur in filtration \mathcal{F}_t .⁴

For $i \in \{1, \dots, N+1\}$ and $\Gamma_i \in \mathcal{E}$, label every event $X_0^{-1}(\Gamma_i) \subset \Omega$ with a unique symbol q , and collect all such symbols in a finite set \mathcal{Q} . Formally, this can be done by means of a surjective function $\Phi : E \rightarrow \mathcal{Q}$, such that $\Phi(g_1) = \Phi(g_2) \iff \exists i \in \{1, \dots, N+1\} : g_1 \in \Gamma_i \ni g_2$.

Pick $(q, q') \in \mathcal{Q} \times \mathcal{Q}$, and for arbitrary $k \in \mathbb{N}$ (recall that X_t is Markov) define the probability measure

$$\mu(q, q') := \mathbb{P} \left[\left\{ \omega \in \Omega \mid \Phi(X_k(\omega)) = q, \Phi(X_{k+1}(\omega)) = q' \right\} \right]. \quad (1)$$

This practically defines a probability distribution that expresses the likelihood of, say, the robot landing at different workspace regions in after time “step” (of duration τ). Indeed, for a given $q \in \mathcal{Q}$, $\mu(q, \cdot)$ is a probability measure on \mathcal{Q} . Let $\mathcal{Q} := 2^{\mathcal{Q}}$ denote the power set of \mathcal{Q} and $\Pi(\mathcal{Q})$ the family of discrete probability spaces⁵ $(\mathcal{Q}, \mathcal{Q}, \mu(q, \cdot))$, parameterized by $q \in \mathcal{Q}$.

Since X is a controlled Markov process, there can be an additional parameterization of the probability spaces $(\mathcal{Q}, \mathcal{Q}, \mu)$ depending on what control law is used to drive the process. For a finite set of possible control laws, each labeled with a symbol $\sigma \in \Sigma$, one obtains a different Markov process for each σ ; the dependence is denoted $X|_{\sigma}$. Subsequently, each σ gives rise to a different family of discrete probability distributions over \mathcal{Q} . Thus given $q \in \mathcal{Q}$, the probability distribution $\mu(q, \cdot)$ changes depending on σ , a fact that is highlighted by writing the left hand side of (1) in the form $\mu_{\sigma}(q, \cdot)$.

This paper adopts a definition of deterministic finite automata that allows only a single state to be initial [31], and builds on it to define a *semiautomaton*. The basic difference between automata and semiautomata is that the latter do not have marked initial and final states. So if automata specify where the system starts from and where it needs to go, semiautomata merely describe how the system can evolve. Given that in the present context this difference is practically inconsequential—the marking of states in automata is invoked in the context of motion planning tasks in Section IV, where there are initial and goal configurations—the discussion that follows switches rather casually between automata and semiautomata.

Definition 1 (Probabilistic semiautomaton; cf. [31]–[33]). A *probabilistic semiautomaton* $\mathbf{A} = (\mathcal{Q}, \Sigma, \Delta)$ is a tuple with components:

\mathcal{Q}	a finite set of states;
Σ	a finite set of input symbols; ^a
Δ	the transition relation. ^b

⁴See Definition 11 in Appendix B.

⁵See Definition 15 in Appendix C.

³See Definition 8 in Appendix B.

^a The input set is sometimes referred to as alphabet, action set or action signature.

^b $\Delta \subseteq \prod_{\sigma \in \Sigma} Q \times \Sigma \times \Pi(Q)$ and when $(q, \sigma, \mu) \in \Delta$ one writes $q \xrightarrow{\sigma} \mu$. The semantics of the later is that for discrete state $q \in Q$, σ determines the probability distribution over the states that may succeed q as X_n evolves. We denote $\mu_\sigma(q, \cdot) := \mu(q, \cdot) |_{(q, \sigma, \mu) \in \Delta}$.

A nondeterministic semiautomaton can be associated to each probabilistic semiautomaton. The difference between the probabilistic semiautomaton A and its associated nondeterministic semiautomaton A_n is that in the latter, the information on transition probabilities is ignored.

Definition 2. Given a probabilistic semiautomaton $A = (Q, \Sigma, \Delta)$, the nondeterministic semiautomaton associated with A is a tuple $A_n = (Q, \Sigma, \Delta_n)$, where the transition relation Δ_n is defined based on Δ , as follows:

$$(q, \sigma, q') \in \Delta_n \iff q' \in \arg \sup \left(\mu(q, \cdot) |_{(q, \sigma, \mu) \in \Delta} \right).$$

In view of the aforementioned mapping of a (continuous time and space) controlled Markov process to a family of finite Markov chains⁶—specifically, one finite Markov chain for a given discrete control input—an uncertain and stochastic physical processes is linked to a probabilistic semiautomaton where the family of discrete probability spaces is not yet fixed.

Definition 3 (Discrete abstraction). For $\sigma \in \Sigma$, consider a finite family of Markov processes $X |_\sigma$ over state space (E, \mathcal{E}) , and let $\{\Gamma_i\}_{i=1}^{N+1}$ be a finite partition of E . With $\Phi : E \rightarrow Q$ being a surjective function and $|Q| = N + 1$, such that $\Phi(g_1) = \Phi(g_2) \iff \exists i \in \{1, \dots, N + 1\} : g_1 \in \Gamma_i \ni g_2$, a probabilistic semiautomaton $A = (Q, \Sigma, \Delta)$ is a discrete abstraction of $X |_\sigma$, if

$$(q, \sigma, \mu) \in \Delta \implies \forall q' \in Q, \mu(q, q') |_{(q, \sigma, \mu) \in \Delta} \equiv \mu_\sigma(q, q') = \mathbb{P} \left[\{ \omega \in \Omega \mid \Phi(X_t |_\sigma) = q, \Phi(X_{t+T} |_\sigma) = q' \} \right]. \quad (2)$$

The last clause of (2) is the formal link between the continuous stochastic process of interest and its discrete probabilistic abstraction. Implication (2) simply states the fact that the transition probabilities in the discrete abstraction of a physical process capture the likelihood of the process to evolve between the corresponding initial and target states. Any probabilistic semiautomaton that can potentially satisfy (2) after a modification of its transition relation Δ will be considered a *model* for the physical process:

Definition 4. Let A be a discrete abstraction of $X |_\sigma$. Suppose A' is a probabilistic semiautomaton whose associated nondeterministic semiautomaton is isomorphic to that of A . Then A' is a model for $X |_\sigma$.

The actual physical process $X |_\sigma$ is thus considered practically unknown; one tries to make predictions about its evolution using a *model* $X' |_\sigma$ in the form of a discrete system that is isomorphic to the process' abstraction.

Consider now two probabilistic semiautomata A, A' with identical set of input symbols Σ , and let R be a relation on $Q \times Q'$. Relation R may be lifted to $\Pi(Q) \times \Pi(Q')$ by means of weights (if the latter exist). Specifically, for $(Q, Q, \mu) \in \Pi(Q)$ and $(Q', Q', \mu') \in \Pi(Q')$, R lifts from Q to Π (cf. [32, Definition 15]), denoted $(Q, Q, \mu) \sim_R (Q', Q', \mu')$, if there exists a function $w : Q \times Q' \rightarrow [0, 1]$ such that

- (a) $\forall (q, q') \in Q \times Q', w(q, q') > 0 \implies (q, q') \in R$;
- (b) $\forall (s, s') \in R$ and $(q, q') \in Q \times Q'$ such that $(s, q) \in \Delta, (s', q') \in \Delta'$, it holds that $\sum_{q' \in Q'} w(q, q') = \mu_\sigma(s, q)$;
- (c) $\forall (s, s') \in R$ and $(q, q') \in Q \times Q'$ such that $(s, q) \in \Delta, (s', q') \in \Delta'$, it holds that $\sum_{q \in Q} w(q, q') = \mu'_\sigma(s', q')$.

Figure 1 depicts an example of a weight function for a pair of related states (s, s') . Assume state s belonging to semiautomaton A , and s' to A' . The weights intuitively express a way of redistributing the probabilities over the states that are successors of s in A , to their related successors states of s' in A' . Note that each pair of related states can have its own weight function. The weight functions essentially serve




Fig. 1: A weight function associated with related states (s, s') . The numbers on the graphs' edges mark the probabilities that the process will evolve along this route under control input a . The dashed arcs link the successors of states s and s' on the same input a , according to the given relation. The weight function assigns positive numbers to each of these links in accordance to conditions (a)–(b)–(c) above.

to as a means to express how the probability distributions in one model can be reshaped in another model when these two models are linked through a (strong) simulation relation:

Definition 5 (Simulation relation; cf. [33]). A relation R on $Q \times Q'$ is a (strong) simulation iff $\forall (q, q') \in R$ and $\sigma \in \Sigma$

$$q' \xrightarrow{\sigma} \mu' \in \Delta' \implies \exists q \xrightarrow{\sigma} \mu \in \Delta \wedge \mu \sim_R \mu'.$$

Then $A = (Q, \Sigma, \Delta)$ strongly simulates $A' = (Q', \Sigma, \Delta')$. Checking whether a relation between two probabilistic (semi)automata is a strong simulation relation is equivalent to solving a maximum flow network problem [34], and has been reduced to a linear program [35]. Indeed, the proof of the next statement is constructed in the spirit of treating the strong simulation verification problem as a linear program.

⁶A Markov chain is a Markov process with a countable state space [30].

Theorem 1. Let $A = (Q, \Sigma, \Delta)$ and $A' = (Q', \Sigma, \Delta')$ be two probabilistic semiautomata whose associated nondeterministic semiautomata are isomorphic.⁷ Let θ denote the isomorphism and assume there exist $s \in Q$ and $\theta(s) = s' \in Q'$ with incoming transitions from all states but no outgoing transitions. Construct a relation $R \subset Q \times Q'$ such that $\forall q \in Q$,

- (i) $(q, \theta(q)) \in R$, and
- (ii) $(q, \theta(s)) \in R$.

If $\forall \sigma \in \Sigma$, the following two conditions hold

- 1) $\mu_\sigma(s) \leq \mu'_\sigma(\theta(s))$, and
- 2) $\mu_\sigma(q) \geq \mu'_\sigma(\theta(q))$, $\forall q \in Q \setminus \{s\}$

then relation R is a strong simulation relation and A strongly simulates A' .

Proof. Fix $\sigma \in \Sigma$. Identify the precursor of $s \in Q$ in A , that is, a $q \in Q$ such that $(q, s) \in \Delta$, and consider all states $\{u_1, \dots, u_{N+1}\}$ reachable from q in A . Without loss of generality, take $s \equiv u_{N+1}$. For $i \in \{1, \dots, N+1\}$, denote $\mu_i = \mu_{\sigma i}$ the probability that q jumps to u_i on σ .

Due to the isomorphism between the associated nondeterministic semiautomata of A and A' , there will be $\theta(q) \in Q'$ and $\{v_1, \dots, v_{N+1}\}$ states in Q' such that for any $i \in \{1, \dots, N+1\}$ we have $(\theta(q), v_i) \in \Delta'$. Once again, assume $s' \equiv v_{N+1}$. Denote μ'_i the probability that $\theta(q)$ transitions to v_i on σ . There are $2N+1$ elements of R pairing one u_i with one v_i : $|\{(u_i, v_i) \in R : (q, u_i) \in \Delta \wedge v_i = \theta(u_i)\}| = 2N+1$. (This is because there are $N+1$ elements due to clause (i) and another N due to (ii).) Associate a weight w_j to each one of those $2N+1$ elements of R , with the convention that for $i \in \{1, \dots, N+1\}$ it is $w(u_i, v_i) = w_i > 0$ and $w(u_i, \theta(u_{N+1})) = w_{N+i} > 0$.

One way to lift R to $\Pi(Q) \times \Pi(Q')$ is to have

$$\mu_i = w_i + w_{N+i} \quad (3a)$$

$$\mu'_i = w_i \quad (3b)$$

for any $i \in \{1, \dots, N+1\}$, and in addition

$$\mu_{N+1} = w_{2N+1} \quad (3c)$$

$$\mu'_{N+1} = \sum_{j=N+1}^{2N+1} w_j. \quad (3d)$$

Denote the vector of weights

$$\mathbf{w} := (w_1, \dots, w_{2N+1})$$

and construct the probability (distribution) vector

$$\mathbf{b} := (\mu'_1, \dots, \mu'_N, \mu_1, \dots, \mu_N, \mu'_{N+1}, \mu_{N+1}).$$

Now the set of equations (3) can be written in matrix form

$$\mathbf{A} \mathbf{w} = \mathbf{b}$$

where matrix \mathbf{A} has the following structure

$$\mathbf{A} := \begin{bmatrix} \mathbf{I}_{N \times N} & \mathbf{0}_{N \times N} & \mathbf{0}_{N \times 1} \\ \mathbf{I}_{N \times N} & \mathbf{I}_{N \times N} & \mathbf{0}_{N \times 1} \\ \mathbf{0}_{1 \times N} & \mathbf{1}_{1 \times N} & 1 \\ \mathbf{0}_{1 \times N} & \mathbf{0}_{1 \times N} & 1 \end{bmatrix},$$

with $\mathbf{I}_{n \times n}$ being the $n \times n$ identity matrix, $\mathbf{0}_{n \times m}$ the $n \times m$ zero matrix, and $\mathbf{0}_{1 \times n}$, $\mathbf{1}_{1 \times n}$ denoting n -dimensional column vectors of zeros and ones, respectively. Note that (3c) and (3b) are decoupled from the other equations; this system reduces to

$$\begin{bmatrix} \mathbf{I}_{N \times N} \\ \mathbf{1}_{1 \times N} \end{bmatrix} \begin{bmatrix} w_{N+1} \\ \vdots \\ w_{2N} \end{bmatrix} = \begin{bmatrix} \mu_1 - \mu'_1 \\ \vdots \\ \mu_N - \mu'_N \\ \mu'_{N+1} - \mu_{N+1} \end{bmatrix}. \quad (4)$$

The last row in (4) is dependent on the ones above: the sum of rows 1 through N gives row $N+1$. Therefore there is always a unique solution to (4):

$$w_{N+1} = \mu_i - \mu'_i \quad w_{2N+1} = \mu'_{N+1} - \mu_{N+1}.$$

However, there are constraints: μ_i and μ'_i are probabilities, and weights w_i cannot be negative. For this solution to be admissible, therefore, the following inequalities should hold:

$$\begin{aligned} \mu_i - \mu'_i &\geq 0 \iff \mu_i \geq \mu'_i \\ \mu'_{N+1} - \mu_{N+1} &\geq 0 \iff \mu_{N+1} \leq \mu'_{N+1}. \end{aligned}$$

Conditions 1) – 2) of the theorem are thus established. \square

To ground Theorem 1 in context, assume that semiautomaton A represents the stochastic dynamics of the physical phenomenon of interest (the transition probabilities of which cannot be known precisely), and that A' corresponds to a hypothesized probabilistic model of that uncertain process. The special states s and $\theta(s) = s'$ play the role of sink or “coffin” states in A and A' respectively; they are states where system trajectories go when the process terminates (e.g., in case of the process attempting but failing to make a particular transition). What Theorem 1 implies is that if the model is constructed in a way that is isomorphic to the actual physics (recall that what is unknown about the physics is the transition function, not the states or input alphabets), then a relation can always be defined in such a way that the physics simulates the model, provided that *the model is more conservative*: the probability of a transition from a state of the model to fail is higher than that the probability in the actual process ($\mu'_{N+1} > \mu_{N+1}$); similarly, the probability of successful transition to another state in the model must be lower compared to the actual physics ($\mu'_i < \mu_i$). When the physics simulates the model, this means that the model is not endowed with more behaviors than those that can be observed in an actual experiment.

Theorem 1 does not provide direct ways of checking or enforcing the validity of the conditions when μ_i (the probability distributions of the physical process) are unknown. If these distributions can be approximated, then it is natural to

⁷See Definition 17 in Appendix C. To particularize to semiautomata, ignore the conditions on the initial and final states.

ask whether simulation relations between discrete stochastic processes can still be established in a probabilistic sense.

The next section aims at making two points. The first is that, indeed, there are principled ways of approximating those distributions, and setting the parameters of the model in such a way so that strong simulation relations between the model and the physical process it represents can be ensured *probabilistically* at any confidence level, given enough *data*. The second point is that strong simulation relations between physics and model are relevant and useful in the application space of planning and navigation for miniature legged robots.

IV. CASE STUDY: MINIATURE LEGGED ROBOTS

Process noise does not always scale down with size, and an example of this can be found in the motion of miniature legged robots. Whereas the uncertainty of ground-robot interaction is noticeable yet manageable in larger meter-scale wheeled and legged robots, for legged robots with form factor in the order of a few inches (Fig. 2a) contact forces have a dramatic effect on motion behavior [28], [29]. The deviation after a few seconds of open-loop motion is comparable to body lengths (Fig. 2b).




Fig. 2: Miniature legged robots are susceptible to noisy ground interactions and exhibit clearly stochastic motion behaviors: (a) A 3D printed biologically-inspired miniature multi-legged robot. (b) When starting from within an initial region open loop robot paths going straight, left, and right start to disperse rapidly.

The goal of this section is to demonstrate how discrete stochastic models in conjunction with simulation relations can be utilized to predict the probability of successful completion of navigation tasks on the robotic platforms of Fig. 2. The particular robot is an open-source 3D-printed miniature hexapod robot design developed at the University of Delaware. The design specifications and CAD files for this robot are freely available at <https://www.thingiverse.com/thing:4084997>.

A. The physics

In this case, the physical process (which in principle includes all robot-ground physical interactions during locomotion) is known only partially, and primarily through experimental observations. For the case of the robot of Fig. 2a, it is assumed that the underlying physics is a Markov process⁸ $X|_{\sigma}$, for $\sigma \in \Sigma := \{l, s, r\}$, where l, s, and r stand for left, straight, and right, respectively. The process has state space (E, \mathcal{E}) , probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and is adapted⁹ to some

filtration $\{\mathcal{F}_t\}$. The transition probability function of $X|_{\sigma}$ is denoted p_{σ} . Both p_{σ} and \mathbb{P} , are in general unknown.

What is given is a body of data $D = \bigcup_{\sigma} D_{\sigma}$, where each D_{σ} is a collection (bundle) of sample paths $X_t(\omega)|_{\sigma}$ all starting in the neighborhood of a common initial state and taking values in E , for $t \in T := [0, \tau] \subset \mathbb{R}_{\geq 0}$ and fixed $\tau > 0$ (see Fig. 2b). Each $X_t(\omega)|_{\sigma}$ is a realization of the process for one of the three available inputs σ . Let $\text{card}(D_{\sigma})$ denote the cardinality of the path bundle D_{σ} , and let $M := \text{card}(D_{\sigma})$. Each path bundle can be seen as an element of the product σ -algebra

$$\mathcal{E}^M = \underbrace{\mathcal{E} \times \dots \times \mathcal{E}}_{M \text{ times}}$$

and the associated product measure \mathbb{P}^M can be defined on $\Omega \times \dots \times \Omega$. Since each sample path is a measurable mapping $\omega \mapsto X_t|_{\sigma}$, one can think similarly of a path bundle as a mapping $\Omega^M \rightarrow E^M$ where $\omega \mapsto D_{\sigma}$ for $\omega \in \Omega^M$.

If the state space E consists of spatial transformations, e.g. rigid body motions, then it will be a Lie group;¹⁰ and if in addition uniformity over E is assumed, for example that the physical space has the same terrain characteristics, constant elevation, etc., then the group action would apply uniformly to all points of the physical workspace of the robot. The significance of these assumptions is that they allow D to contain a single bundle of sample paths $X_t|_{\sigma}$, one for each $\sigma \in \Sigma$; one would not need a separate data set for each potential initial condition. Without loss of generality, it is assumed that the common starting neighborhood of all observed sample paths surrounds the group's identity element.

The average of all sample paths in bundle D_{σ} is a curve denoted \bar{X}_{σ} (examples of which are the thick green curves in Fig. 2b). Formally, $\bar{X}_{\sigma} : T \rightarrow E$ is a map for which

$$\bar{X}_{\sigma}(s) := \mathbb{E}[\pi_s(X_t|_{\sigma})], \quad s \in T, \quad X_t|_{\sigma} \in D_{\sigma}.$$

With a slight abuse of notation, we use just $X_s|_{\sigma}$ to denote $\pi_s(X_t|_{\sigma})$, until we switch to canonical processes.

Each \bar{X}_{σ} serves as the representative of bundle D_{σ} , and can be regarded in itself as a spatial transformation taking the Lie group's identity element to the average of $X_t|_{\sigma}$ for all sample paths in bundle D_{σ} . These representatives will henceforth be referred to as *motion primitives*.

B. The abstraction of the physics

Based on the available motion primitives \bar{X}_{σ} , and given an initial condition, a discrete reachability tree can be constructed for the unknown stochastic process X_t . In particular, the average paths \bar{X}_{σ} can be concatenated according to the group action: for $\sigma_1, \sigma_2 \in \Sigma$, if one assumes that motion primitives \bar{X}_{σ_1} and \bar{X}_{σ_2} have end points $g_{\sigma_1}, g_{\sigma_2} \in E$, respectively, then the concatenation $\bar{X}_{\sigma_2} \bar{X}_{\sigma_1}$ is understood as a spatial transformation that maps the identity to $g_{\sigma_2} g_{\sigma_1}$. If the selected initial condition is some $g_0 \in E$, then the result of executing the motion primitive \bar{X}_{σ_2} after \bar{X}_{σ_1} is a rigid body transformation to $g_{\sigma_2} g_{\sigma_1} g_0$.

⁸See Definition 13 in Appendix B.

⁹See Definition 12 in Appendix B.

¹⁰In the present case study, $E = \text{SE}(2)$ —see Appendix D.

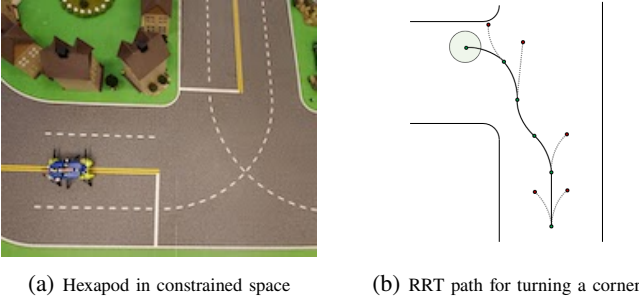


Fig. 3: Planning motion primitive-based paths for the hexapod of Fig. 2a. (a) The hexapod is placed at an intersection of a scaled-down model of a city [36], with the intention of running in open loop a sequence of motion primitives in order to make a left turn. (b) The primitive-based randomly-exploring random tree (RRT) constructed for this mission with the selected path to a goal disk-shaped region.

To make this idea more concrete, consider a case where a robot is operating in a constrained planar environment, where the obstacle-free space has uniform characteristics. The goal is for the robot to reach a region $F \subset E$, starting from an initial condition that is distributed around $g_0 \in E$ according to density P_0 . There are multiple ways of planning that motion; one that is convenient for the purposes of the analysis here is the use of a primitive-based RRT [37]. This reachability tree has nodes where g_0 is mapped under concatenations of motion primitives in Σ (Fig. 3).

Let the nodes of the reachability tree be N in number. The indexing convention is that node 1 is the root of the tree (at g_0); node N is the single leaf node at g_{N-1} , which is inside the target region. With the addition of a sink node, these N nodes “seed” $E = \text{SE}(2)$ (see Section III) and form the finite set \mathcal{Q} with $|\mathcal{Q}| = N + 1$. The collection of disjoint compact neighborhoods $\{\Gamma_i\}_{i=1}^N$, together with $\Gamma_{N+1} := \text{SE}(2) \setminus \bigcup_{i=1}^N \Gamma_i$, partition $\text{SE}(2)$ into $N + 1$ blocks. Neighborhood Γ_1 is a ϖ -content tolerance region¹¹ around g_0 ; assuming that distribution P_0 of initial conditions is known, this region can be constructed to contain exactly the ϖ percentile of initial conditions at a γ confidence level. Neighborhood Γ_N is the desired target area for the process in E . The remaining neighborhoods will be similarly defined as tolerance regions; the following section gives more details.

Definition 6. Let a (symbolically) controlled continuous Markov process $X|_\sigma$ with state space (E, \mathcal{E}) over a finite set of control symbols $\sigma \in \Sigma$ be characterized by a transition probability function $p_\sigma(\text{d}y, t + \text{d}t | x, t) = P_{t+\text{d}t}(x, \text{d}y; \sigma)$ with known initial distribution P_0 . Then a discrete abstraction of $X|_\sigma$ is a probabilistic automaton $A = (\mathcal{Q}, q_1, \Sigma, \Delta, \mathcal{Q}_F)$, with components described as follows:

- \mathcal{Q} the set of states;^a
- $q_1 \in \mathcal{Q}$ the initial state;^b
- Σ the finite set of input symbols;^c
- Δ the transition relation;^d
- $q_N \in \mathcal{Q}$ the final state.^e

¹¹See Definition 14 in Appendix B.

- ^a The N nodes of the reachability tree g_i , and the elements of $\mathcal{Q} \setminus \{q_{N+1}\}$, are linked through a bijection $E \rightarrow \mathcal{Q} \setminus \{q_{N+1}\}$, where $g_i \mapsto \Phi(g_i) = q_{i+1}$. Each $q_i \in \mathcal{Q} \setminus \{q_{N+1}\}$ is a representative of a compact neighborhood $\Gamma_i \subset E$ around a node. State q_{N+1} is a representative of $\Gamma_{N+1} = E \setminus \bigcup_{i=1}^N \Gamma_i$ and is a “sink” state in the sense that there is no outgoing transition from that state.
- ^b The initial state q_1 , is associated with the root of the reachability tree at $g_0 \in E$.
- ^c The input set contains a symbol for each one of the control inputs (e.g., the motion primitives).
- ^d $\Delta \subseteq \mathcal{Q} \times \Sigma \times \Pi(\mathcal{Q})$ where $(q_i, \sigma, \mu) \in \Delta$ implies that for any $q_j \in \mathcal{Q} \setminus \{q_{N+1}\}$, and for a fixed $\tau \in (0, T]$,

$$\begin{aligned} \mu &= \mu_\sigma(q_i, q_j) = \int_E P_\tau(g, \text{d}g') \mathbb{1}_{\Gamma_j}(g') \\ &= \int_{g \in \Gamma_i} \int_{g' \in \Gamma_j} p(X_\tau |_\sigma = g' | X_0 |_\sigma = g) \text{d}g' \text{d}g . \end{aligned} \quad (5a)$$

For A it holds that

$$\begin{aligned} \exists q_k \in \mathcal{Q} \setminus \{q_{N+1}\} : \\ \int_{g \in \Gamma_i} \int_{g' \in \Gamma_k} p(X_\tau |_\sigma = g' | X_0 |_\sigma = g) \text{d}g' \text{d}g \\ > \int_{g \in \Gamma_i} \int_{g' \in \Gamma_j} p(X_\tau |_\sigma = g' | X_0 |_\sigma = g) \text{d}g' \text{d}g \\ \implies \mu_\sigma(q_i, q_j) := 0 . \end{aligned} \quad (5b)$$

It is assumed that for every $q_i \in \mathcal{Q}$ and $\sigma \in \Sigma$, it is $(q_i, \sigma, q_{N+1}) \in \Delta$, and

$$\mu_\sigma(q_i, q_{N+1}) = 1 - \sum_{k=1}^N \mu_\sigma(q_i, q_k) . \quad (5c)$$

- ^e The final state is associated with the leaf of the reachability tree that is inside F , that is, $q_N = \Phi(g_{N-1})$.

From (5) it follows that for a given $\sigma \in \Sigma$ and $q_i \in \mathcal{Q} \setminus \{q_{N+1}\}$, there exist a transition from q_i to some $q_j \in \mathcal{Q} \setminus \{q_{N+1}\}$ only if $\mu_\sigma(q_i, q_j)$ is maximal over all pairs (q_i, q_j) . Transitions to other states are considered unintentional, associated with failure, and are redirected to the sink state q_{N+1} .

Probabilistic automaton A is a finite representation of the stochastic process at hand and it abstracts it *accurately*, albeit this abstraction not *concrete*: distributions μ_σ are unknown, and with the exception of first and last neighborhoods, i.e. Γ_1 and Γ_N , all other regions have not been specified. The next section describes how to make the model concrete.

C. The model for the abstraction

Given an (uncertain) discrete abstraction A of a finite family of continuous Markov processes $X|_\sigma$, and a body of data D consisting of observations of the underlying physical process, a finite *empirical model* A' can be constructed. The

discrete abstraction A and the finite model A' are two probabilistic automata with isomorphic associated nondeterministic automata.¹²

Definition 7. Consider the probabilistic automaton $A = (Q, q_1, \Sigma, \Delta, Q_f)$ of Definition 6. A model for A is a probabilistic automaton $A' = (Q, q_1, \Sigma, \Delta', Q_f)$, which is isomorphic to A and whose transition relation Δ' satisfies the following three conditions: for every $\sigma \in \Sigma$,

- 1) $(q, \sigma, \mu') \in \Delta' \iff (q, \sigma, \mu) \in \Delta$
- 2) $\forall j \neq N+1, \mu_\sigma(q_i, q_j) > 0 \Rightarrow 0 < \mu'_\sigma(q_i, q_j) < \mu_\sigma(q_i, q_j)$
- 3) $\mu_\sigma(q_i, q_{N+1}) < \mu'_\sigma(q_i, q_{N+1})$.

It follows directly from Theorem 1 that A' is strongly simulated by A . Motion plans synthesized based on A' can be executed on A with a probability that is at least as high as their probability in A' . Utilizing a procedure for Probabilistic Model Validation (PMV) [38], the probability distributions of A' can be determined from observations of $X|_\sigma$ so that A' is simulated by A . The inclusion of the behavior of A' into that of A guarantees the assertion above.

Given the controlled continuous Markov process $(\Omega, \mathcal{F}, \mathbb{P}; \{X_t|_\sigma : t \in T\})$ capturing the physics, take $X'|_\sigma$ as the *canonical* [30] controlled continuous Markov process $(E, \mathcal{E}, \mathbb{P}'; \pi_t : t \in T)$ with $\mathbb{P}' = \mathbb{P} \circ X_t^{-1}$, satisfying

$$\mathbb{E}[X'_t|_\sigma] = \bar{X}_\sigma \quad t \in T. \quad (6)$$

Fixing $\sigma \in \Sigma$, let measure \mathbb{P}'_σ belong to a family of probability measures \mathcal{P}'_σ (e.g., ones giving rise to a family of Gaussian processes with mean \bar{X}_σ) and be parameterized by a vector ξ_σ of model parameters taking values in some bounded set Ξ_σ so that $\bigcup_{\xi \in \Xi} \mathbb{P}'_\sigma(\xi_\sigma) = \mathcal{P}'_\sigma$. Thus the law of process X'_t , namely $\mathbb{P}' = \mathbb{P}' \circ (X'_t)^{-1}$, depends implicitly on ξ_σ . Picking a measure \mathbb{P}_{Ξ_σ} on Ξ_σ induces a measure on \mathcal{P}'_σ and produces a *distribution of processes* $X'|_\sigma$. The idea now is to sample \mathbb{P}_{Ξ_σ} and use data D in order to identify probably approximate near optimal fits for ξ_σ so that A' is strongly simulated by A with probability $1 - \rho$ at confidence level γ .

The PMV-based parameter selection method reported here does not depend on the choice of \mathbb{P}_{Ξ_σ} nor on the specifics of $X'|_\sigma$. (See Example 1 for a particular implementation of $X'|_\sigma$.) Without yet particularizing $X'|_\sigma$, and having defined Γ_1 and Γ_N (Section IV-B), one completes the definition of the neighborhood regions by referencing the given reachability tree (e.g. Fig. 3) and identifying the sequences of motion primitives $\bar{X}_{\sigma_1}, \dots, \bar{X}_{\sigma_k}$ that connect the tree's root to any node g_i other than the final.

Reachability graph being a tree, implies that every node g_i has a unique parent denoted $g_{\text{parent}(i)}$ from which it can be reached via some specific \bar{X}_σ . Now for $i = 1, \dots, N$, each Γ_i is defined as the element of \mathcal{E} coinciding with the ϖ -probability content (a, φ) tolerance rectangle of \mathbb{P}'_σ , when $g_{\text{parent}(i)}$ is assumed uniformly distributed in $\Gamma_{\text{parent}(i)}$. In other words, if j denotes $\text{parent}(i)$, then Γ_i will be the smallest subset

of E for which

$$\mu'_\sigma(q_i, q_j) := \int_{g \in \Gamma_j} \int_{g' \in \Gamma_i} p'(X'_\tau|_\sigma = g' | X'_0|_\sigma = g) dg' dg \geq \varpi. \quad (7)$$

Note that the left hand side of (7) depends on ξ_σ due to \mathbb{P}'_σ , as the latter determines the distribution of $g' = X'_\tau|_\sigma g$ over E . If necessary, one can potentially pick different ϖ for different σ , making the probability content of each tolerance region explicitly dependent on σ and write it as $\varpi(\sigma)$. Selecting $\varpi(\sigma)$ sufficiently small ensures that all Γ_i neighborhoods are disjoint.

Ideally μ'_σ is independent of (q_i, q_j) ; that is, execution of σ results in the same transition probabilities to target and sink states uniformly on A' . Such a property facilitates the construction of A' as well as the estimation of reachability probabilities over different paths on A' . This happens if

$$(X'_\tau|_\sigma)^{-1}(\Gamma_i) = (X'_\tau|_\sigma)^{-1}(\Gamma_j) \quad (8)$$

for all $i, j \in \{1, \dots, N\}$ and $\sigma \in \Sigma$. Section IV-E demonstrates how (8) can be enforced for the particular case study at hand through a particular choice of the underlying probability space $(\Omega', \mathcal{F}', \mathbb{P}')$. Such choices of $(\Omega', \mathcal{F}', \mathbb{P}')$ are likely to require an explicit dependence of ϖ on σ .

The *center* of each neighborhood Γ_i is the associated reachability tree node g_{i-1} . For an arbitrary pair (q_i, q_j) giving $\mu_\sigma(q_i, q_j) \neq 0$, center $q_j = g_{j-1}$ will naturally coincide with $\mathbb{E}[\mathbb{P} \circ (X'_\tau|_\sigma)^{-1}]$ when $X'_0|_\sigma = q_i = g_{i-1}$.

Example 1. For the planar motion with the motion primitives \bar{X}_σ of Fig. 2b for $\sigma \in \{l, r, s\}$, g_i can be assumed an element of $\text{SE}(2)$ (Fig. 5) and \bar{X}_σ is in the role of a planar rigid body transformation. Note that the translation component of these transformations can be parameterized in the form of circular arcs with primitive-specific radii, in terms of a , the arc length, and φ , the absolute orientation of the frame which the circle containing the arc is tangent to at the origin (Fig. 4). The rotation component can still be expressed in terms of a rotation angle θ , around the z axis. In that light, we take (a, φ, θ) to be the (new) canonical parameterization of E .

The uniqueness of this parameterization is that under the assumptions about how noise is injected in the motion of the system (see Appendix A), disturbances in θ enter with those in φ in affine form, greatly facilitating the analytical propagation of uncertainty. For example, as illustrated in Fig. 4, if one assumes that a frame g moves along a circular arc (thick curve) of length a and of rotation φ relative to the horizontal tangent, it will end up at the end point of the arc, (x, y) pointing in the tangent to that point direction, which is equal to 2ζ . In this way, translation and rotation appear coupled, but the third (rotational) degree of freedom is restored by considering an independent rotation in place at (x, y) which will give $g \in \text{SE}(2)$ its final orientation, θ .

For a given primitive motion σ , take the three variables $(a_\sigma, \varphi_\sigma, \theta_\sigma)$ parameterizing E , to be independent, normally distributed random variables parameterized by $\sigma \in \Sigma$, with

¹²See Definition 4. For the isomorphism, refer to Definition 17 in Appendix C.

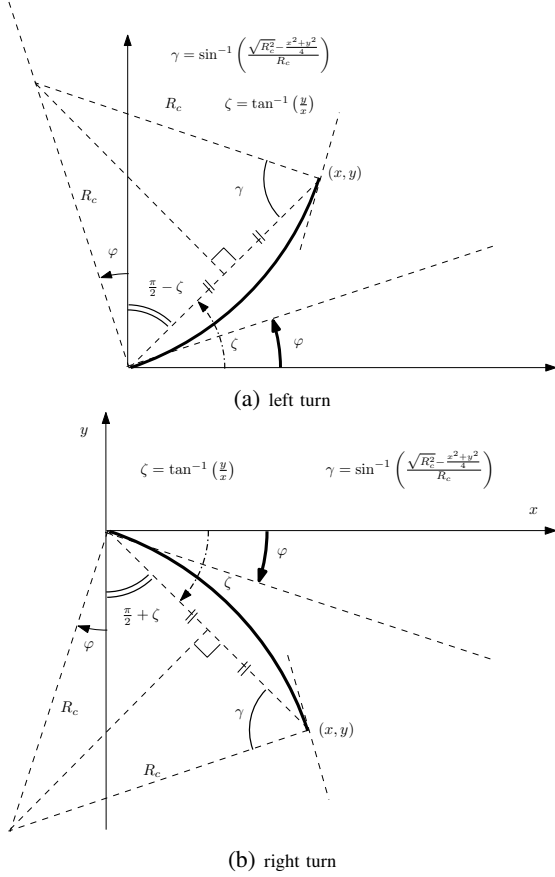


Fig. 4: Parameterization of primitives in terms of arc length and rotation, and explicit relations between Cartesian coordinates (x, y) and proposed parameterization (a, φ) : (a) Left turn. (b) Right turn.

means $\bar{X}_\sigma(\tau)$ and variances ξ_σ . A benefit of this parameterization under the assumptions made is that θ_σ rotations in transformation g_{i-1} simply add to φ_σ to provide the initial (orientation) conditions for the subsequent transformation g_i . Once the effect of θ_σ is incorporated in this way, the pose distribution propagation analysis can proceed in two dimensions in between primitives.

Since each $\Gamma_i \subset E$ is associated with a specific σ , it can be defined in this case as the rectangular $(a_\sigma, \varphi_\sigma, \theta_\sigma)$ -tolerance region translated to $q_i = g_{i-1}$ (see Fig. 5a for an illustration of the tolerance region in the $(a_\sigma, \varphi_\sigma)$ subspace). When one changes the position parameterization in E from (a, φ) to the Cartesian (x, y) , the rectangular tolerance regions deform and appear as in Fig. 5b, which depicts a transition from a neighborhood Γ_1 to a neighborhood Γ_2 along a left-turn primitive \bar{X}_1 . Starting from $X'_0|_1 = g_0 = q_1$, the end location $X'_\tau|_1$ is distributed about $g_1 = q_2$ according to $\mathbb{P}' \circ (X'_\tau)^{-1}$ —a binormal distribution in the (a, φ) plane (Fig. 5a) appearing as a banana distribution (cf. [39]) in Fig. 5b.

The banana-like distribution of Fig. 5b is in fact the convolution of $\mathbb{P}' \circ (X'_\tau)^{-1}$ with a uniform distribution \mathcal{U} over Γ_1 . Specifically, for the initial conditions of the primitive $X'_\tau|_1$ it is assumed that $X'_0 \sim \mathcal{U}(\Gamma_1)$. Justification for this choice is given next in Section IV-D as it discusses the process for tuning the parameter vector ξ for the model automaton in order to be

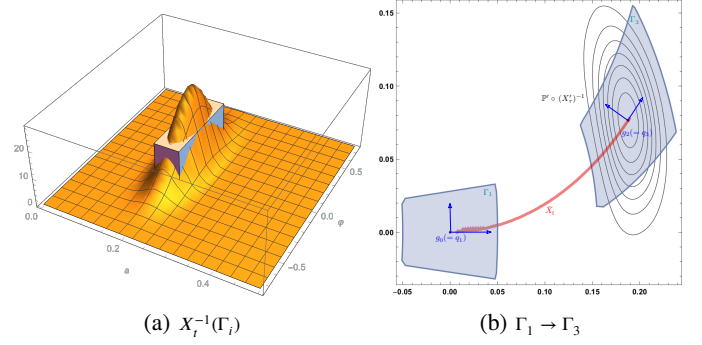


Fig. 5: The effect of workspace parameterization on motion primitive distributions: (a) In the a - φ space, events (outcomes of a (a, φ) sampling experiment) can be independent and normally distributed, and the Γ_i regions can be defined as rectangular tolerance regions; (b) A Cartesian parameterization deforms the Γ_i regions—here a transition from region Γ_1 to region Γ_3 via a left-turn primitive is a random experiment resulting in a distribution $P'_\tau(g, dg') = \mathbb{P}' \circ (X'_\tau|_1)^{-1}$ of poses around $\bar{X}_1 g_0$ that resembles the “banana” distributions of [39].

strongly simulated by the discrete abstraction A .

D. Model parameter selection

The macroscopic kinematics of bio-inspired multi-legged robots like the one in Fig. 2a can be modeled using a one-parameter mechanical abstraction called the Switching Four-bar Mechanism (SFM) [40]. (Details about this kinematic model, and a new stochastic extension of it, can be found in Appendix A.) In this paper, the cartesian position configuration variables of the SFM abstraction are re-parameterized in terms of a and φ as depicted in Fig. 4. Under the assumption that random disturbances are injected at leg touchdown, and perturb the system normally over (arc) translation a and pivot rotation φ relative to the initial pose, the end-point uncertainty after the execution of a motion primitive \bar{X}_σ is captured by $\mathbb{P}' \circ (X'_\tau|_\sigma)^{-1}$. The latter is parameterized in terms of ξ , and the main goal of this section is to set ξ to some ξ_σ in order to match the variability observed in experimental observations.

In general, given σ and ξ , and depending on the sample ω' drawn from Ω' , a sample path $X'_\tau(\omega'; \xi)|_\sigma$ originating in some Γ_i has a certain probability of hitting Γ_j , given by $\mu'_\sigma(q_i, q_j)$ in (7). This probability depends on the choice of model parameter ξ , which controls the dispersion of sample paths about \bar{X}_σ . Since measure \mathbb{P}'_σ of the model process is known, $\mu'_\sigma(q_i, q_j)$ can, in principle, be computed accurately. For the process $X|_\sigma$ describing the actual physics, however, the probability of hitting Γ_j from Γ_i is unknown. Under the assumptions of workspace uniformity introduced in Section IV-B, the statistics of transition of $X|_\sigma$ from Γ_i to Γ_j is captured by that of transitioning from Γ_1 to $\bar{X}_\sigma \Gamma_1$ upon σ . Analogously, $X|_\sigma$ missing its target neighborhood is an event

$$\Omega_\sigma := \left\{ \omega \in \Omega : X_0(\omega) \in \Gamma_1, X_\tau(\omega)|_\sigma \notin \bar{X}_\sigma \Gamma_0 \right\} \in \mathcal{F}$$

that has probability $\rho_\sigma := \mathbb{P}(\Omega_\sigma)$. Given that what is desired is for the physics X to simulate the model X' , this probability places an upper bound on the probability of events in which experimental data (observations of $X|_\sigma$) disagree with model

predictions (based on $X' |_\sigma$). In other words, the model should not allow paths to miss their target neighborhood on a σ transition with probability more than ρ_σ . The challenge is that \mathbb{P} is unknown and thus ρ_σ cannot be computed; it may, however, be empirically approximated based on data.

In order for A to simulate A' , Theorem 1 implies that model transitions from some Γ_i to some other Γ_j should have probability of success at most as those of the actual physics. The theorem's condition on transitions to discrete states other than the sink is rewritten in terms of μ'_σ and ρ_σ as

$$\rho_\sigma < 1 - \mu'_\sigma(q_i, q_j) \quad \forall \sigma \in \Sigma, \forall q_i, q_j \in Q \setminus \{q_{N+1}\} \quad (9)$$

meaning that the probability of the model failing to transition on σ from q_i to q_j is larger than that of the actual physics, which is ρ_σ . Recall that μ'_σ is implicitly a function of ξ , and so it can be adjusted to satisfy (9). This adjustment needs to work for all $(q_i, q_j) \in Q$, and therefore the explicit dependence of μ'_σ on this pair will be dropped for notation convenience. Rewriting (9), define

$$G(\xi) := \rho_\sigma + \mu'_\sigma(\xi) < 1 \quad (10)$$

and set

$$\xi^* := \arg \sup_{\xi \in \Xi} \mu'_\sigma(\xi) \quad \text{subject to } \rho_\sigma < 1 - \mu'_\sigma(\xi) \quad (11)$$

as the solution that maximizes the left hand side of (10). Given that ρ_σ is unknown, one now asks how small should μ'_σ be made through the choice of ξ^* to satisfy the above inequality without being overly conservative? An answer to this question can be provided by randomized algorithms.

Let Γ_σ denote the ϖ -probability content tolerance (a, φ) -rectangle of \mathbb{P}'_σ when $X_0 |_\sigma \in \Gamma_0$. Then the empirical average of the probability of violation on the data bundle D_σ

$$\hat{\rho}_\sigma := \frac{1}{M} \sum_{X_t |_\sigma \in D_\sigma} \left[1 - \mathbb{1}_{\Gamma_\sigma}(X_\tau |_\sigma) \right]$$

will be a random variable that is expected to tend to the actual ρ_σ as $M \rightarrow \infty$. For a small accuracy parameter ε and at confidence level γ , the Chernoff bound [41]

$$M \geq \frac{1}{2\varepsilon^2} \ln \frac{2}{\gamma} \quad (12)$$

suggests conservatively just how big M should be in order for $\hat{\rho}$ to be ε -close to ρ with probability $1 - \gamma$.

If a closed-form expression of μ'_σ as a function of ξ is available, then the selection of ξ may be straightforward. More often than not, however, the relationship is not explicit; then, randomized algorithms can once again be used to select an appropriate set of model parameters ξ in the spirit of (12).

Specifically, consider a finite sample $\{\xi_\ell\}_{\ell=1}^{M'} \subset \Xi$, satisfying $\mathbb{E}[\xi_\ell] = \xi_\sigma$. One may now be tempted to relax the search over the whole Ξ in (11) into a search for ξ^* over this *finite* collection of ξ_ℓ by requiring somewhat more conservatively that $\forall D_\sigma \in \mathcal{E}^M$

$$\hat{\mu}'_\sigma := \max_{1 \leq \ell \leq M'} \mu'_\sigma(\xi_\ell) < 1 - (\hat{\rho}_\sigma(D_\sigma) + \varepsilon) . \quad (13)$$

Yet this will come at a price: one has to accept that there will be a set of model parameters, the effect of which is not

going to be represented adequately by any member of the finite collection $\{\xi_\ell\}_{\ell=1}^{M'}$. This subset of Ξ is assumed to have measure α , and ideally this α is comfortably small. Also note that the last quantifier is important: it signifies the implicit requirement that the model parameter ξ_ℓ^* which realizes the maximum $\hat{\mu}'_\sigma$ should work not only for the bundle D_σ provided as “training” data, but for *any* path bundle.

Parameterizing the model with a value taken from $\{\xi_\ell\}_{\ell=1}^{M'}$, denote $\Gamma_\sigma(\xi_\ell)$ the γ -confidence and ϖ -probability content tolerance region of $\mathbb{P}'_\sigma(\xi_\ell)$ when $X_0 |_\sigma$ is assumed distributed *uniformly* in Γ_1 . Then for

$$\hat{\rho}_{\sigma\ell} := \frac{1}{M} \sum_{X_t |_\sigma \in D_\sigma} \left[1 - \mathbb{1}_{\Gamma_\sigma(\xi_\ell)}(X_\tau |_\sigma) \right] \quad (14)$$

one can reasonably claim that $\hat{\rho}_{\sigma\ell} \geq \hat{\rho}_\sigma$ as long as the variance of P_0 is smaller than that of $\mathcal{U}(\Gamma_1)$. Keeping γ and ϖ constant for all neighborhoods $\{\Gamma_i\}_{i=1}^N$ allows to generalize the statement $\hat{\rho}_{\sigma\ell} \geq \hat{\rho}_\sigma$ to all such neighborhoods and nodes of the reachability tree, and consequently, states $q \in Q \setminus \{q_{N+1}\}$ of probabilistic automaton A' . Then the condition for A to simulate A' is satisfied by construction. What is more, defining the model's transition probabilities as in (13) ensures that A' is not overly conservative while being simulated by A . Thus, a specification for selecting the model parameters ξ_ℓ can be formulated, by combining (13) with (14) in a way that elucidates how the simulation relation condition depends on sample sizes M and M' and the choice of ξ_ℓ :

$$\begin{aligned} \hat{\mu}'_\sigma + \hat{\rho}_{\sigma\ell} < 1 - \varepsilon &\iff \max_{1 \leq \ell \leq M'} \mu'_\sigma(\xi_\ell) \\ &- \frac{1}{M} \sum_{X_t |_\sigma \in D_\sigma} \mathbb{1}_{\Gamma_\sigma(\xi_\ell)}(X_\tau |_\sigma) + 1 < 1 - \varepsilon . \end{aligned} \quad (15)$$

A randomized algorithm now maximizes the left hand side, making the above inequality condition as tight as reasonably (with respect to sample sizes) possible.

Corollary 1. *With confidence $1 - \gamma$, the quantity*

$$\hat{G} := \max_{1 \leq \ell \leq M'} \left\{ \mu'_\sigma(\xi_\ell) - \frac{1}{M} \sum_{X_t |_\sigma \in D_\sigma} \mathbb{1}_{\bar{X}_\sigma \Gamma_\sigma(\xi_\ell)}(X_\tau |_\sigma) \right\} + 1 \quad (16)$$

is a (Type 3) probably approximate near maximum of $G(\xi)$ to accuracy ε and level α if

$$M' \geq \frac{\ln \frac{2}{\gamma}}{\ln \frac{1}{1-\alpha}} \quad \text{and} \quad M \geq \frac{1}{2\varepsilon^2} \ln \frac{4M'}{\gamma} .$$

Proof. Direct application of [42, Algorithm 1]. \square

Proposition 1. *If \hat{G} is a probably approximate near maximum of $G(\xi)$ to accuracy ε , level α , and confidence γ , then the probability of A simulating A' is at least $1 - \alpha$ with that same confidence.*

Proof. With confidence $1 - \gamma$ and at level α one writes $G(\xi) > \hat{G} + \varepsilon$, which in turns leads to

$$\begin{aligned} G(\xi) > \hat{G} + \varepsilon &\implies G(\xi^*) > \hat{G} + \varepsilon \iff \rho_\sigma + \mu'_\sigma(\xi^*) \\ &> \max_{1 \leq \ell \leq M'} \mu'_\sigma(\xi_\ell) + 1 - \frac{1}{M} \sum_{X_t \in D_\sigma} \mathbb{1}_{\tilde{X}_\sigma \Gamma_\sigma(\xi_\ell)}(X_t |_\sigma) + \varepsilon \\ &= \hat{\mu}'_\sigma + \hat{\rho}_\sigma + \varepsilon. \end{aligned}$$

By definition [42] of a (Type 3) maximum, $\Pr \left\{ \xi \in \Xi : G(\xi^*) > \hat{G} + \varepsilon \right\} \leq \alpha$, which is identical to $\Pr \left\{ \xi \in \Xi : \rho_\sigma + \mu'_\sigma(\xi) > \hat{\mu}'_\sigma + \hat{\rho}_\sigma + \varepsilon \right\} \leq \alpha$, implying that

$$\Pr \left\{ \xi \in \Xi : G(\xi) > \hat{G} + \varepsilon \right\} \leq \alpha.$$

With confidence $1 - \gamma$, therefore, there is $1 - \alpha$ probability to have $\hat{\mu}'_\sigma + \hat{\rho}_\sigma + \varepsilon \geq \rho_\sigma + \mu'_\sigma(\xi)$, and given that $1 > \hat{G} + \varepsilon$ (see (15)) it follows that (9)—which is the condition for the simulation relation—holds with probability $1 - \alpha$. \square

With this mathematical machinery in place, a value of 0.3 for the probability of violation ρ_σ is arbitrarily selected as acceptable—since the same probability is adopted for all primitives, the subscript is henceforth dropped. The objective now is to select ξ_σ so that the empirical estimate $\hat{\rho}_\sigma$ is no more than ε above the target, with confidence γ and at level α . Parameters ε , γ , and α are (arbitrarily, again) set at 0.25, 0.1, and 0.25, respectively, and with that, Corollary 1 and Proposition 1 would guarantee a 75% (i.e. $1 - \alpha$) chance that the robot physics A will simulate the probabilistic automaton model A' (a statement made with 60% (i.e. $1 - \rho$) confidence, given the accepted probability of violation).

To this end, a data set D composed of three data bundles D_σ of size $M = 460$ each is collected with the robot of Fig. 2a. Each bundle is associated with the robot performing straight, left turn, and right turn maneuvers starting from a common, fixed initial location, and are labeled s, l, and r, for straight, left turn, and right turn, respectively, yielding $\Sigma = \{s, l, r\}$. The means $(\bar{a}_\sigma, \bar{\varphi}_\sigma)$ associated with each D_σ are calculated and used to tune in a least-squares way the SFM mean displacement parameters [40] so that $\tilde{X}_\sigma = \mathbb{E}[\pi_s(X_t |_\sigma)]$. Although (6) cannot be guaranteed to be satisfied *exactly*, one can still ensure that $\mathbb{E}[X_t |_\sigma]$ and $\mathbb{E}[X'_t |_\sigma]$ are not statistically different.

From this point, a PMV procedure is initiated to find ξ_σ^* in the form of standard deviations $(\tilde{a}_\sigma, \tilde{\varphi}_\sigma)$, for which *no more than 30% ($\rho = 0.3$) of sample paths in any experimental data bundle may fall outside a 90% inclusion tolerance region Γ around \tilde{X}_σ* . This randomized optimization process yields the values reported in Table I, where the data bundle standard deviations $(\tilde{a}_\sigma, \tilde{\varphi}_\sigma)$ are listed as well for comparison (results can vary given the randomized nature of the optimization):

The distributions $\mathbb{P}' \circ (X'_t |_\sigma)^{-1}$, parameterized with the values in the left two columns of Table I, after convolved with a uniform distribution over the region of origin yield tolerance regions Γ_i of the form similar to that of Fig. 5b. Indeed, starting uniformly from a rectangular a - φ tolerance region Γ_0 with 90% inclusion centered around the starting configuration, and with the ξ_σ^* values of Table I, the data bundles D_σ result in containment percentages in Γ_2 , Γ_4 , and Γ_3 regions (each

σ	\tilde{a}'_σ	ξ_σ^* $\tilde{\varphi}'_\sigma$	$\tilde{\theta}'_\sigma$	\tilde{a}_σ	D_σ $\tilde{\varphi}_\sigma$	$\tilde{\theta}_\sigma$
s	0.01608	0.1827	0.3616	0.01573	0.1005	0.3616
l	0.01400	0.1980	0.2593	0.01440	0.2170	0.2593
r	0.01232	0.1460	0.2545	0.01359	0.1613	0.2545

TABLE I: PMV results for the model primitive standard deviations ξ_σ^* (left) juxtaposed with the data bundle D_σ statistics (right).

corresponding to motion primitives s, l, and r), which are 83.91%, 81.21% and 78.82%, respectively. In comparison, the probability that the model $X'_t |_\sigma$ predicts for landing in Γ_i for $i \in \{2, 3, 4\}$, with initial configurations *uniformly distributed* in Γ_1 (see Fig. 2b) but *without noise in the θ_σ dimension*, is $\mu'_\sigma(\xi_\sigma^*) = 0.6344$, $\forall \sigma \in \Sigma$. When the uncertainty in θ_σ is incorporated (which needs to happen whenever a motion primitive is concatenated with another) these probabilities drop to $\mu'_s(\xi_s^*) = 0.2730$, $\mu'_l(\xi_l^*) = 0.3197$, and $\mu'_r(\xi_r^*) = 0.3100$ (see Fig. 7). The computation of μ'_σ , which generally involves the integration of non-standard numerically approximated probability distributions, is straightforward is not necessarily trivial. On a 2.4 GHz Quad-Core Intel Core i5 processor with 16GB of memory, it can take a total of up to 35 minutes in a Mathematica implementation. It is noteworthy, however, that a computation such as the one detailed above is performed off line, and *only once per each control law $\sigma \in \Sigma$* . For once μ'_σ is calculated, and under assumptions similar in nature to those made in the description of Example 1 and the beginning of Section IV-A, any future instance of application of σ directly inherits the results of that earlier computation.




Fig. 6: The tolerance regions reachable from a starting configuration, and the distribution of experimentally produced primitive end-points (red crosses) executed from the same initial robot pose. The black dots on the boundary of the regions mark the vertices of polygonal numerical approximations of the tolerance regions, while the dots in the interior show the position of $X'_t |_\sigma$. The green arcs from the initial pose on the left to each $X'_t |_\sigma$ endpoint illustrate \tilde{X}_σ .

Once all μ'_σ are computed, estimating the probability of transitioning from state q to state q' becomes as straightforward as multiplying the transition probabilities along the path that connects q to q' (see Fig. 7).

E. Experimental validation

Recall the robot navigation scenario introduced in Section IV-B. The robot is a six-legged bio-inspired crawler

shown in more detail in Fig. 2a, and the experimental testbed environment is shown in Fig. 3a. The motion planning GOAL area is marked with a disk in Fig. 3b, and reaching it requires the robot to make a sharp left turn while staying in the road boundaries of the setup in Fig. 3a. Figure 3b shows a primitive-based RRT path to the goal region.

A rectangular tolerance region for this joint distribution is mapped in the Euclidean (robot work)space E through primitive \tilde{X}_σ as a distorted trapezoid like that shown in Fig. 6. The size of this tolerance region is naturally dependent on the variances of the a and φ distributions assumed for the model. These variances will be the model parameters selected through the optimization process outlined in Section IV-D, and the key parameter ultimately determining the size of the regions is the probability of violation ρ_σ . For the purpose of this case study, tolerance intervals with 90% containment are used.

The three primitives (in the form of paths $\tilde{X}_\sigma \subset E$) can now be used by an RRT planner to connect nodes in a tree joining START to a neighborhood within GOAL; a typical output of such a planner can be seen in Fig. 3b. Based on this planner output the process model A' is constructed, yielding the probabilistic automaton shown in Fig. 7.




Fig. 7: The graph of probabilistic automaton A' , which serves as a model for the abstraction of the stochastic process that generated the sample paths in Fig. 6 when it executes sequences of motion primitives realizable on the RRT tree of Fig. 3b. Node 12 is the sink.

Trials Bundle #	size	Primitive sequence traversal success probabilities				
		s	sl	slr	slrl	slrll
MAXIMUM THEORETICALLY ALLOWED/EXPERIMENTALLY OBSERVED						
1	105	1.00/0.88	0.55/0.30	0.18/0.09	0.03/0.00	—
2	199	1.00/0.70	0.45/0.24	0.13/0.06	0.04/0.02	0.01/0.00
3	159	1.00/0.78	0.57/0.27	0.24/0.13	0.08/0.08	0.01/0.00
4	95	1.00/0.69	0.51/0.29	0.21/0.12	0.08/0.08	0.02/0.00
model	—	0.63	0.20	0.06	0.02	0.01

TABLE II: Comparison between the probabilities of successively hitting Γ_i along an RRT path between model predictions (last row) and four experimentally observed path bundles.

A total of 558 experimental trials of the RRT solution of Fig. 3 are executed, grouped into four data bundles as shown on the two left columns of Table II. The remaining columns of Table II index the different prefixes of the primitive sequence

from initial to goal configuration and give two pieces of information about the samples in each data bundle: the bottom number is the empirical (observed) probability with which the paths in the data bundle hit the corresponding Γ_i region starting from the initial configuration; the top number is the theoretically maximum probability of success on hitting that region, given that PMV allows for 30% probability of violation, i.e., three in ten experimental observations not to agree with the model. The latter figures were obtained by multiplying the transition probabilities found in the path that is associated with each control sequence in Fig. 7. The data of Table II confirm (a) that the model predictions on the probability of the robot successively hitting the different target regions on the RRT tree follow the trend observed in the experimental data and (b) that the model predictions remain consistently conservative compared to experimental observations, at least until the tail of the cumulative hitting distribution where the data samples become prohibitively small to provide meaningful statistics.

V. CONCLUSION

For a partially known physical process that exhibits stochastic behavior, it is possible to construct an abstraction in the form of a probabilistic automaton. The parameters of the probability distributions for the transitions in this automaton can be tuned based on experimental observations, in a way that this abstract model is strongly simulated by the physical process—even if a concrete model for the actual physics is not available. Such an abstract model is more conservative than the actual physics in terms of the probability of transitioning between states, an attribute that allows the designer to have confidence that control strategies synthesized based on the model will be at least as probable of succeeding in the real system as they are predicted to do in the abstract model.

NOTATION USED

Δ	transition relation
Γ	compact neighborhood of some point
γ	confidence level
$\mathbb{1}_\Gamma$	indicator function of Γ
\mathbb{P}	probability measure
\mathcal{E}	σ -algebra on E
\mathcal{E}^T	σ -algebra on E^T
\mathcal{F}	σ -algebra
\mathcal{F}_t	filtration
\mathcal{P}'	class of distributions for the continuous model
e_i	basis vector of the Lie algebra on SE(2)
\mathfrak{g}	element of the Lie algebra on SE(2)
A	probabilistic semiautomaton
A_n	nondeterministic semiautomaton
Ω	sample space
ω	sample point
Ω^M	product sample space
Φ	morphism for $\{\Gamma_i\}$
ϕ^{lo}	leg lift-off angle
ϕ^{td}	leg touchdown angle
Π	set of discrete probability spaces
π_i	valuation map

ρ_σ	probability of violation for executing input σ
Σ	finite input symbol set (alphabet)
σ	input alphabet symbol
\sim_R	relation on distributions
ϖ	percentile of population in tolerance region
A	probabilistic automaton
D_σ	bundle of sample paths from same initial point
E	range of a stochastic process
E^T	space of functions from T to E
F	target region
G	randomized optimization cost function
g	crawler pose: element of the Lie group $SE(2)$
g_i	reachability tree node and neighborhood Γ_{i+1} center
M	cardinality of the sample path bundle D_σ
p	probability law of stochastic process
P_0	distribution of initial conditions
P_t	transition semigroup
Q	finite set of discrete states
R	relation between discrete states
W	stochastic process jump rate
X	continuous stochastic process
$X _\sigma$	stochastic process controlled by σ

APPENDIX

A. The stochastic SFM model

For miniature legged robots like the one shown in Fig 2a that move quasi-statically (e.g., at relatively low speeds and excluding flight phases), a particular kinematic model, called the SFM [40] (see Fig. 8), can serve as a low-dimensional kinematic abstraction of the actual physics of the motion. The particular abstraction has proven particularly convenient in the sense that it affords *closed-form* (Fig. 8b) analytic descriptions of state propagation (i.e. for yaw $\Delta\theta$ and geometric center displacement GG'). There is experimental evidence that this model can reasonably describe the quasi-static motion a variety of different miniature crawling robots [43].

If the end-points of the virtual legs of the model—for instance, those marked with red dots in Fig. 8a—are assumed to be hinged on the ground and the legs rotate about those hinges, then the mechanism will propel itself forward. The rotation of the model's (virtual) legs, which generate motion for the model, is parameterized in terms of touchdown ϕ_s^{td} (initial) and lift-off ϕ_s^{lo} (final) angles. Touchdown angles define the configuration of the four-bar mechanism at the beginning of a step, while lift-off angles mark the configuration where the active group of legs (solid thick line segments adjacent to contact points) are losing contact with the ground (at the red points), and the other pair of legs which was previously in flight (dashed line segments) now reacquiring contact (at the green dots). Depending on whether the lift-off angles are chosen to be opposite of the touchdown angles or not, the mechanism can be made to displace its geometric center along a straight line segment or a circular arc (Fig. 8b), motivating the macroscopic robot path representation of Fig. 4.

Let the pose of the SFM model be $g(\phi_s) \in SE(2)$, where $s = 1$ when AO_1-BO_2 (Fig. 8b) are on the ground and ϕ_1 is

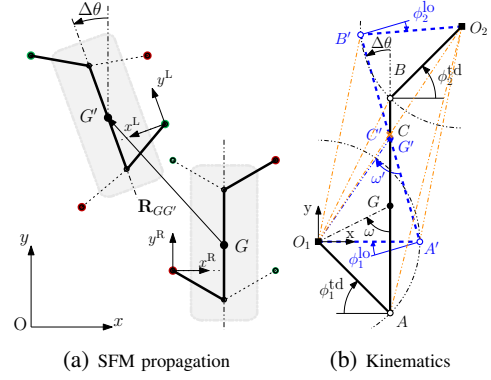


Fig. 8: To produce bio-inspired locomotion gaits, the legs of each side and across the left and right sides of a multi-robot platform have to be coupled, either mechanically or electronically. The SFM kinematic model groups the coupled legs on each side of the legged platform that are in contact with the ground, and abstracts them into a single virtual leg per side. In this form, the ground, legs, and torso of the robot form a closed kinematic chain which is modeled by a single degree-of-freedom planar four-bar mechanism parameterized by one of the virtual legs angle relative to the robot's body [38].

assumed driving the mechanism, and $s = 3$ otherwise. The pose evolution *over a single step* can be parameterized in terms of ϕ_s instead of time t ; the function $\phi_s(t)$ can bring time back into the picture, if necessary. The limits of the range of ϕ_s during a stride, i.e., the associated touchdown and lift-off angles can be selected through an optimization algorithm [40] so that the evolution of the model matches—e.g., in a least-squares sense—to the experimentally observed averages \bar{X}_σ . Assuming therefore that an appropriate choice of geometric model parameters ensures $\mathbb{E}[X_\tau | \sigma] = g(\phi_s^{\text{lo}})$, this process yields a deterministic model that evolves on the same reachability graph captured by automaton A that represents the model of the physics (e.g., the RRT of Fig. 3b).

Under appropriate assumptions on how noise can perturb the evolution of this deterministic model, a particular stochastic process can arise: (a) stochastic uncertainty is injected in the motion of the system only at *leg touchdown* instances, while the stance phase remains noise-free and deterministic; and (b) injected noise characteristics are Gaussian. Under these two assumptions, the pose of the SFM model, g , is distributed according to a density function p' , the evolution of which obeys a forward differential Chapman–Kolmogorov equation of a particular form [5]. This particular form dictates that $g(\phi_s)$ is a random variable that evolves *deterministically* during a step, and *jumps randomly* when leg pair AO_3-BO_4 succeed pair AO_1-BO_2 (or vice-versa) in ground contact, i.e., when $\phi_s = \phi_s^{\text{lo}}$. This particular form of the forward Chapman–Kolmogorov has no diffusion term, and the transition rate is

$$W(g' | g, \phi_s) = \lim_{\delta \rightarrow 0} \frac{1}{\delta} p'(g', \phi_s + \delta | g, \phi_s) \\ =: \begin{cases} 0 & \phi_s \neq \phi_s^{\text{lo}} \\ \frac{1}{\sqrt{2\pi}\zeta} \exp\left(-\frac{\|g-g'\|^2}{2\zeta^2}\right) & \phi_s = \phi_s^{\text{lo}} \end{cases},$$

where ζ is a positive constant parameter and $\|\cdot\|$ is an

appropriate metric on SE(2). With that, the propagation of the model's pose during a single step becomes a continuous stochastic process, characterized by a forward differential Chapman–Kolmogorov equation on SE(2) where the transition probability function p' evolves according to [44]

$$\begin{aligned} \frac{\partial p'(g, \phi_s | g_0, \phi_s^{\text{td}})}{\partial \phi_s} &= \int \left[W(g | g', \phi_s^{\text{td}}) p'(g', \phi_s | g_0, \phi_s^{\text{td}}) \right. \\ &\quad \left. - W(g' | g, \phi_s) p'(g, \phi_s | g_0, \phi_s^{\text{td}}) \right] dg' \\ &- \frac{d}{dt} \left[p'(g \circ \exp(te_1)) p'(g \circ \exp(te_2)) p'(g \circ \exp(te_3)) \right] \Bigg|_{t=0} \\ &\quad \cdot \left(g^{-1} \frac{\partial g}{\partial \phi_s} \right)^{\vee}. \quad (17) \end{aligned}$$

This equation characterizes a process with sample paths having drift and jumps (cf. [5, Fig. 3.2(b)] and note that here the jumps actually occur at regular intervals). There is one equation of the form (17) when legs AO_1 – BO_2 are on the ground, and another when AO_3 – BO_4 are on the ground. The stochastic model dynamics exhibit switching between the two equations when leg lift-off occurs.

Quantization of touchdown and lift-off angle configurations can lead to three basic motion primitives: straight motion, left turn, and right turn [40]. The model parameters are tuned [38] for each primitive so that the kinematic behavior matches the experimentally observed robot motion (Fig. 2b) when the same primitives are realized on the physical robot. The drift component of the platform's dynamics alone, as expressed by the second term on the right-hand-side of (17), would produce deterministic motion for the SFM.

When the model executes a (single) motion primitive σ with initial condition in Γ_i for some $i \in \{1, \dots, N\}$, it has some finite probability of landing in any one $\Gamma_j \in \{\Gamma_1, \dots, \Gamma_{N+1}\}$ for $j \neq i$. (Failure to transition outside the original neighborhood after executing the primitive is considered a failure, and by default it maps to Γ_{N+1} instead.) The probability of transitioning from Γ_i to Γ_j , denoted $\mu'_\sigma(q_i, q_j)$, is now encoded in (17) and quantified as follows:

$$\mu'_\sigma(q_i, q_j) = \int_{g_i \in \Gamma_i} \int_{g_j \in \Gamma_j} p'(g, \phi_s^{\text{td}} | g_0, \phi_s^{\text{td}}) dg_j dg_i. \quad (18)$$

B. Stochastic processes

Definition 8 (σ -algebra [30]). Let S be a set. A collection \mathcal{F} of subsets of S is called an algebra on S if it is closed under the finitely many set operations:

$$\begin{aligned} S &\in \mathcal{F} \\ B \in \mathcal{F} &\implies B^c := S \setminus B \in \mathcal{F} \\ B, C \in \mathcal{F} &\implies B \cup C \in \mathcal{F}. \end{aligned}$$

\mathcal{F} is a σ -algebra on S if for any countable family of sets $\{B_n\}$ with $B_n \in \mathcal{F}$ for any $n \in \mathbb{N}$, it holds $\bigcup_n B_n \in \mathcal{F}$. If \mathcal{B} is a class of subsets of S , then the σ -algebra generated by \mathcal{B} , denoted $\sigma(\mathcal{B})$, is the smallest σ -algebra on S that contains \mathcal{B} . \diamond

Definition 9 (Probability space cf. [32]). A probability space is a triplet $(\Omega, \mathcal{F}, \mathbb{P})$ where Ω is a set (typically called sample space), \mathcal{F} is a σ -algebra on Ω , and \mathbb{P} is a function from \mathcal{F} to $[0, 1]$ such that $\mathbb{P}[\Omega] = 1$ and for any collection $\{C_i\}_i$ of at most countably many disjoint elements of \mathcal{F} , it is $\mathbb{P}[\bigcup_i C_i] = \sum_i \mathbb{P}[C_i]$ (\mathbb{P} is a probability measure). \diamond

Definition 10 (Stochastic process [30]). A stochastic process with time-parameter set T , state-space (E, \mathcal{E}) and probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is a collection of (E, \mathcal{E}) -valued random variables.

Definition 11 (Filtration [30]). A filtration on probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is an increasing family $\{\mathcal{F}_n : n \in \mathbb{N}\}$ satisfying $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_\infty := \sigma(\bigcup_n \mathcal{F}_n) \subseteq \mathcal{F}$. \diamond

Definition 12 (Adapted process [30]). A process $X = \{X_n : n \in \mathbb{Z}\}$ carried by $(\Omega, \mathcal{F}, \mathbb{P})$ is said to be adapted (to the given filtration) if, for every $n \in \mathbb{Z}$, X_n is \mathcal{F}_n -measurable.

Definition 13 (Markov process, cf. [30]). Let $\{P_t\}$ denote a family of kernels $P_t : E \times \mathcal{E} \rightarrow [0, 1]$ such that for any bounded measurable function f on E , it holds:

- 1) for $t \geq 0$ and $x \in E$, $P_t(x, \cdot)$ is a measure on \mathcal{E} with $P_t(x, E) \leq 1$;
- 2) for $t \geq 0$ and $\Gamma \in \mathcal{E}$, $P_t(\cdot, \Gamma)$ is \mathcal{E} -measurable;
- 3) for $t \geq 0$, $x \in E$, and $\Gamma \in \mathcal{E}$,

$$P_t f(x) := (P_t f)(x) = \int_E P_t(x, dy) f(y).$$

A Markov process X with state space $(E, \mathcal{E}) \subseteq (\mathbb{R}^n, 2^{\mathbb{R}^n})$, is an E -valued stochastic process

$$\left(\Omega, \{\mathcal{F}_t : t \geq 0\}, \{X_t : t \geq 0\}, \{P_t : t \geq 0\}, \{\mathbb{P}^x : t \geq 0\} \right)$$

adapted to filtration $\{\mathcal{F}_t\}$, such that for $0 \leq s \leq t$, $x \in E$, $\mathbb{E}^x[f(X_{s+t}) | \mathcal{F}_s] = (P_t f)(X_s)$, \mathbb{P}^x a.s. \diamond

Definition 14 (cf. [45]). Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and let X be a Markov process with state space (E, \mathcal{E}) , for $t \in [0, \tau] = T$, and with initial condition $X_0 = x$. For process X , and for $M \in \mathbb{N}$, a ϖ -content tolerance region $S(X(\omega_1), \dots, X(\omega_M))$ at confidence $1 - \gamma \in (0, 1)$ is a statistic over $(\mathcal{E}^T)^M$ taking values in \mathcal{F} , for which the probability that $\mathbb{P}^x((X_t)^{-1}(S)) \geq \varpi$ is at least $1 - \gamma$. \diamond

In other words, one has confidence $1 - \gamma$ that region S contains at least ϖ percentile of the population of the sample paths of X that start from $x \in E$ when \mathbb{P} is the measure on Ω . Usually the parameters of the probability measure \mathbb{P} are unknown. For example, in cases where \mathbb{P} is Gaussian [46], [47], the mean, or the variance, or both, can be unknown; different estimates of the region are obtained for each case.

C. Finite automata

Definition 15 (Discrete probability space [32]). A probability space $(\Omega, \mathcal{F}, \mathbb{P})$ ¹³ is discrete if $\mathcal{F} = 2^\Omega$ and for all $C \subseteq \Omega$, $\mathbb{P}[C] = \sum_{x \in C} \mathbb{P}[\{x\}]$. \diamond

¹³See Definition 9 in Appendix B.

It follows that in a discrete probability space there can only be countably many elements with positive probability measure [32].

Definition 16 (Semiautomaton; cf. [31]). *A semiautomaton (SA) is a triple $A = (Q, \Sigma, T)$ where Q is a finite set of states, Σ is a finite alphabet, $T : Q \times \Sigma \rightarrow 2^Q$ is the transition map which can be expanded recursively in the usual way.* \diamond

In general, T maps to a set, in which case we say that A is *non-deterministic in transitions*. We write $T(q, \sigma) \downarrow$ if $T(q, \sigma) \neq \emptyset$, and $T(q, \sigma) \uparrow$ otherwise. A *run* in A on some word $w = \sigma^{(0)}\sigma^{(1)} \dots \in \Sigma^*$, is a finite sequence of states $\rho = q^{(0)}q^{(1)}q^{(2)} \dots \in Q^*$ such that $q^{(i+1)} \in T(q^{(i)}, \sigma^{(i)})$. In this case we say that ρ is *generated by w* .

The transition map in A can be made *total* by adding a state in Q called *sink*, such that for all other states $q \in Q$ for which there exists a symbol $\sigma \in \Sigma$ that cannot trigger a transition from q , namely $T(q, \sigma) \uparrow$, we augment the transition relation with $T(q, \sigma) = \text{sink}$. By definition, we set $T(\text{sink}, \sigma) = \text{sink}$, for all $\sigma \in \Sigma$. Total semiautomata are those in which the transition map is total.

Definition 17 (Automata isomorphism [31]). *Let $A = (Q, q_0, \Sigma, \Delta, Q_F)$ and $A' = (Q', q'_0, \Sigma, \Delta', Q'_F)$ be two automata with the same input alphabet Σ . An isomorphism $\theta : Q \rightarrow Q'$ is a function satisfying*

- 1) $\theta(q_0) = q'_0$;
- 2) $q \in Q_F \implies \theta(q) \in Q'_F$;
- 3) $\forall q \in Q, \forall \sigma \in \Sigma, \theta(\Delta(q, \sigma)) = \Delta'(\theta(q), \sigma)$. \diamond

This definition is naturally extended to semiautomata by dropping conditions 1) and 2).

D. Planar motions

The Euclidean motion group $\text{SE}(2)$ represents rigid-body motions on the plane. It consists of elements of the form

$$g = \begin{bmatrix} \cos \theta & \sin \theta & x \\ -\sin \theta & \cos \theta & y \\ 0 & 0 & 1 \end{bmatrix}.$$

The group operation is denoted \circ and for $\text{SE}(2)$ it is essentially matrix multiplication. Spatial velocities of rigid-body motions are expressed in the form

$$\mathfrak{g} := g^{-1} \dot{g} = \begin{bmatrix} 0 & \dot{\theta} \cos \theta - \dot{y} \sin \theta \\ -\dot{\theta} & \dot{x} \sin \theta + \dot{y} \cos \theta \\ 0 & 0 & 0 \end{bmatrix}$$

and form the Lie algebra $\mathfrak{se}(2)$ associated with $\text{SE}(2)$. Both $\text{SE}(2)$ and $\mathfrak{se}(2)$ are of dimension 3, and are parameterized by x , y , and θ . In this context, the exponential map $\exp : \mathfrak{se}(2) \rightarrow \text{SE}(2)$ becomes the typical matrix exponential.

The Lie algebra $\mathfrak{se}(n)$ is locally identified with \mathbb{R}^n using the map $^\vee : \mathfrak{se}(n) \rightarrow \mathbb{R}^n$, e.g.,

$$\mathfrak{g} \mapsto \begin{pmatrix} \dot{\theta} \\ \dot{x} \\ \dot{y} \end{pmatrix} := (g^{-1} \dot{g})^\vee.$$

Similarly, the map $^\wedge$ for $n = 2$ takes elements of \mathbb{R}^3 to $\mathfrak{se}(2)$:

$$\begin{pmatrix} \dot{\theta} \\ \dot{x} \\ \dot{y} \end{pmatrix}^\wedge := \begin{bmatrix} 0 & \dot{\theta} & \dot{x} \\ -\dot{\theta} & 0 & \dot{y} \\ 0 & 0 & 0 \end{bmatrix}.$$

The basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ for $\mathfrak{se}(2)$ can be generated by applying the $^\wedge$ operator on the nominal basis of \mathbb{R}^3 :

$$\mathbf{e}_1 = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \mathbf{e}_2 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \mathbf{e}_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

For an analytic function $f : \text{SE}(2) \rightarrow \mathbb{R}$, the right Lie derivative along $\mathfrak{g} \in \mathfrak{se}(2)$ is expressed as

$$(\tilde{\mathfrak{g}}^r f)(g) = \frac{d}{dt} \left\{ f(g \circ \exp(t\mathfrak{g})) \right\} \Big|_{t=0}.$$


REFERENCES

- [1] X. Kan, J. Thomas, H. Teng, H. G. Tanner, V. Kumar, and K. Karydis, "Analysis of ground effect for small-scale UAVs in forward flight," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3860–3867, 2019.
- [2] X. Yu, M. A. Hsieh, C. Wei, and H. G. Tanner, "Synchronous rendezvous for networks of marine robots in large scale ocean monitoring," *Frontiers in Robotics and AI*, vol. 6, no. 76, p. doi.org/10.3389/frobt.2019.00076, 2019.
- [3] S. Das, E. B. Steager, M. A. Hsieh, K. J. Stebe, and V. Kumar, "Experiments and open-loop control of multiple catalytic microrobots," *Journal of Micro-Bio Robotics*, vol. 14, pp. 25–34, 2018.
- [4] A. Julius and G. Pappas, "Approximations of stochastic hybrid systems," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1193–1203, 2009.
- [5] C. Gardiner, *Stochastic Methods; A Handbook for the Natural and Social Sciences*, 4th ed. Springer, 2009.
- [6] B. Oksendal, *Stochastic Differential Equations: An Introduction with Applications*, 6th ed. Springer-Verlag Berlin Heidelberg, 2003.
- [7] X. Mao, *Stochastic Differential Equations and Applications*, 2nd ed. Woodhead Publishing Ltd, 2011.
- [8] J. Aguilar and D. I. Goldman, "Robophysical study of jumping dynamics on granular media," *Nature Physics*, vol. 12, pp. 278–283, 2016.
- [9] S. Shen, N. Michael, and V. Kumar, "Autonomous multi-floor indoor navigation with a computationally constrained MAV," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2011, pp. 20–25.
- [10] A. Hosoi and D. I. Goldman, "Beneath our feet: Strategies for locomotion in granular media," *Annual Reviews in Fluid Mechanics*, vol. 47, pp. 231–453, 2015.
- [11] G. S. Chirikjian, *Stochastic Models, Information Theory, and Lie Groups*. Birkhäuser, 2009, vol. 1: Classical Results and Geometric Methods.
- [12] O. Maler, A. Pnueli, and J. Sifakis, "On the synthesis of discrete controllers for timed systems," in *Symposium on Theoretical Aspects of Computer Science*, E. Mayr and C. Puech, Eds., vol. 900, 1995, pp. 229–242.
- [13] M. M. Jr., A. Davitian, and P. Tabuada, "PESSOA: A tool for embedded control software synthesis," in *Computer Aided Verification*, T. Touili, B. Cook, and P. Jackson, Eds., vol. 6174. Springer-Verlag, 2010, pp. 566–569.
- [14] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 53, no. 1, p. 28700297, 2008.
- [15] R. Alur, T. Kenzinger, G. Lafferriere, and G. Pappas, "Discrete abstractions of hybrid systems," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 971–984, 2000.
- [16] G. Pola, A. Girard, and P. Tabuada, "Approximately bisimilar symbolic models for nonlinear control systems," *Automatica*, vol. 44, no. 10, pp. 2508–2516, 2008.
- [17] G. Pola and P. Tabuada, "Symbolic models for nonlinear control systems: alternating approximate bisimulations," *SIAM Journal of Control and Optimization*, vol. 48, no. 2, pp. 719–733, 2009.
- [18] G. Reibig, "Computing abstractions of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 56, no. 11, pp. 2583–2598, 2011.
- [19] A. Abate, J.-P. Katoen, and A. Mereacre, "Quantitative automata model checking of autonomous stochastic hybrid systems," in *Hybrid Systems: Computation and Control*, M. Caccamo, E. Frazzoli, and R. Grosu, Eds. ACM, 2011, pp. 83–92.
- [20] J. Sproston, "Discrete-time verification and control for probabilistic rectangular hybrid automata," in *Proceedings of the IEEE International Conference on Quantitative Evaluation of Systems*, 2011, pp. 79–88.

- [21] S. Esmail, Z. Soudjani, and A. Abate, "Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes," *SIAM Journal of Applied Dynamical Systems*, vol. 12, no. 2, pp. 921–956, 2013.
- [22] A. Abate, J.-P. Katoen, J. Lygeros, and M. Pradini, "Approximate model checking of stochastic hybrid systems," *European Journal of Control*, vol. 6, pp. 624–641, 2010.
- [23] M. Lahijanian, S. Anderson, and C. Belta, "A probabilistic approach for control of a stochastic system from LTL specifications," in *Proceedings of the IEEE Conference on Decision and Control*, 2009, pp. 2236–2241.
- [24] J. Sproston, "Exact and approximate abstractions for classes of stochastic hybrid systems," in *Proceedings of the 14th Workshop on Automated Verification of Critical Systems*, 2014, pp. 109–123.
- [25] M. Zamani, P. Esfahani, R. Majumdar, A. Abate, and J. Lygeros, "Bisimilar finite abstractions of stochastic control systems," in *Proceedings of the IEEE Conference on Decision and Control*, 2013, pp. 3926–3931.
- [26] M. Zamani, I. Tkachev, and A. Abate, "Bisimilar symbolic models for stochastic control systems without state-space discretization," in *Hybrid Systems: Computation and Control*, M. Fränzle and J. Lygeros, Eds. ACM, 2014, pp. 41–50.
- [27] M. Zamani, P. Esfahani, A. Abate, and J. Lygeros, "Symbolic models for stochastic control systems without stability assumptions," in *Proceedings of the European Control Conference*, 2013, pp. 4257–4262.
- [28] X. A. Wu, T. M. Huh, A. Sabin, S. A. Suresh, and M. R. Cutkosky, "Tactile sensing and terrain-based gait control for small legged robots," *IEEE Transactions on Robotics*, vol. 36, no. 1, pp. 15–27, 2020.
- [29] K. Karydis, I. Poulakakis, and H. G. Tanner, "A navigation and control strategy for miniature legged robots," *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 214–219, 2017.
- [30] L. Rogers and D. Williams, *Diffusions, Markov Processes and Martingales; Volume 1: Foundations*, 2nd ed. Cambridge University Press, 2000.
- [31] M. V. Lawson, *Finite Automata*. Chapman & Hall/CRC, 2004.
- [32] R. Segala and N. Lynch, "Probabilistic simulations for probabilistic processes," *Nordic Journal of Computing*, vol. 2, no. 2, pp. 250–273, 1995.
- [33] M. Stoelinga, "An introduction to probabilistic automata," *Bulletin of the European Association for Theoretical Computer Science*, no. 78, pp. 176–198, 2002.
- [34] C. Baier, B. Engelen, and M. Majster-Cederbaum, "Deciding bisimilarity and similarity for probabilistic processes," *Journal of Computer and Systems Sciences*, vol. 60, no. 1, pp. 187–231, 2000.
- [35] L. Zhang and H. Hermanns, *Automated Technology for Verification and Analysis: Proceedings of the 5th International Symposium*. Springer Berlin Heidelberg, 2007, vol. 4762, ch. Deciding simulations on probabilistic automata, pp. 207–222.
- [36] A. Stager, L. Bhan, A. Malikopoulos, and L. Zhao, "A scaled smart city for experimental validation of connected and automated vehicles," in *15th IFAC Symposium on Control in Transportation Systems*, 2018, pp. 130–135.
- [37] K. Karydis, Y. Liu, I. Poulakakis, and H. G. Tanner, "Navigation of miniature legged robots using a new template," in *Proceedings of the IEEE Mediterranean Conference on Control and Automation*, 2015, pp. 1112–1117.
- [38] K. Karydis, I. Poulakakis, and H. G. Tanner, "Probabilistically valid stochastic extensions of deterministic models with uncertainty," *The International Journal of Robotics Research*, vol. 34, no. 10, pp. 1278–1295, 2015.
- [39] A. Long, K. Wolfe, M. Mashner, and G. Chirikjian, "The banana distribution is Gaussian: a localization study with exponential coordinates," in *Robotics: Science and Systems VIII*, N. Roy, P. Neuman, and S. Shrivastha, Eds., 2013, pp. 265–272.
- [40] K. Karydis, Y. Liu, I. Poulakakis, and H. G. Tanner, "A template candidate for miniature legged robots in quasi-static motion," *Autonomous Robots*, vol. 38, pp. 193–209, 2015.
- [41] M. Vidyasagar, "Statistical learning theory and randomized algorithms for control," *IEEE Control Systems Magazine*, pp. 69–85, December 1998.
- [42] —, "Randomized Algorithms for Robust Controller Synthesis using Statistical Learning Theory," *Automatica*, vol. 37, no. 10, pp. 1515–1528, 2001.
- [43] K. Karydis, A. Stager, H. G. Tanner, and I. Poulakakis, "Experimental validation of a template for navigation of miniature legged robots," in *Experimental Robotics: The 15th International Symposium on Experimental Robotics*, O. Khatib, D. Kulic, Y. Nakamura, and G. Venture, Eds. Springer, 2017 (in print).
- [44] W. Park, Y. Wang, and G. S. Chirikjian, "The path-of-probability algorithm for steering and feedback control of flexible needles," *The International Journal of Robotics Research*, vol. 29, no. 7, pp. 813–830, 2010.
- [45] D. A. S. Fraser and I. Guttman, "Tolerance regions," *The Annals of Mathematical Statistics*, vol. 27, no. 1, pp. 162–179, 1956.
- [46] M. Siotani, "Tolerance regions for a multivariate normal distribution," Stanford University, Technical Report 5, 1963.
- [47] V. Witkovsky, "On the exact tolerance intervals for univariate normal distribution," in *Proceedings of the 10th International Conference on Computer Data Analysis and Modeling: Theoretical and Applied Stochastics*, S. Aivazian, P. Filzmoser, and Y. Kharin, Eds., Minsk, Belarus, 2013, pp. 130–137.



Bert Tanner (SM'08) received his Ph.D. in mechanical engineering from the NTUA, Athens, Greece, in 2001. He was a postdoctoral researcher at the University of Pennsylvania from 2001 to 2003, and subsequently took a position as an assistant professor at the University of New Mexico. In 2008 he joined the Department of Mechanical Engineering at the University of Delaware, where he is currently a professor, and director of the Center for Autonomous and Robotics Systems. Tanner received NSF's Career award in 2005. He is a fellow of the ASME, and a senior member of IEEE. He has served in the editorial boards of the IEEE Transactions on Automatic Control, the IEEE Robotics and Automation Magazine and the IEEE Transactions on Automation Science and Engineering. He is currently the Specialty Chief Editor for Frontiers in Robotics and AI Multi-Robot Systems, and an associate editor for Automatica, as well as for Nonlinear Analysis Hybrid Systems. He has also been serving in several conference editorial boards of both IEEE Control Systems and IEEE Robotics and Automation Societies.



Adam Stager was born in New Jersey, USA, in 1988. He received a B.S. in Mechanical Engineering from the University of Delaware, Newark, DE, USA in 2011, then worked in manufacturing as a project engineer for several years before returning to the University of Delaware for graduate school. In 2020, he earned his Ph.D degree in mechanical engineering and subsequently worked as a postdoctoral researcher in the Department of Plant and Soil Sciences at the University of Delaware. He founded, and is currently the CEO of TRIC Robotics LLC, a robotics company which is specializing in agricultural technology using automated systems to replace chemical pesticides on farms.