

# Probabilistic Model Checking of Robots Deployed in Extreme Environments

Xingyu Zhao,<sup>1</sup> Valentin Robu,<sup>1,2</sup> David Flynn,<sup>1</sup> Fateme Dinmohammadi,<sup>1</sup>  
Michael Fisher,<sup>3</sup> Matt Webster<sup>3</sup>

<sup>1</sup>School of Engineering & Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, U.K.

<sup>2</sup>Center for Collective Intelligence, Massachusetts Institute of Technology, Cambridge MA 02139, U.S.

<sup>3</sup>Department of Computer Science, University of Liverpool, Liverpool L69 3BX, U.K.

Email: {xingyu.zhao, v.robust, d.flynn, f.dinmohammadi}@hw.ac.uk, {mfisher, matt}@liverpool.ac.uk

## Abstract

Robots are increasingly used to carry out critical missions in extreme environments that are hazardous for humans. This requires a high degree of operational autonomy under uncertain conditions, and poses new challenges for assuring the robot's safety and reliability. In this paper, we develop a framework for probabilistic model checking on a layered Markov model to verify the safety and reliability requirements of such robots, both at pre-mission stage and during runtime. Two novel estimators based on conservative Bayesian inference and imprecise probability model with sets of priors are introduced to learn the unknown transition parameters from operational data. We demonstrate our approach using data from a real-world deployment of unmanned underwater vehicles in extreme environments.

## 1 Introduction

Extreme environments, as a term used by UK EPSRC<sup>1</sup> to denote environments that are remote and hazardous for humans, are one of the most promising application areas in which robots can be deployed to carry out a task, such as inspecting oil and gas equipment, maintaining offshore wind turbines or monitoring nuclear reactors. Interaction with human operators is often infeasible in such remote environments, so autonomy becomes a key requirement, meaning robots have to learn to adapt when performing tasks in changing and unexpected circumstances (Lane et al. 2016).

However, unforeseen behaviour of a robot could result either in failure of its own mission or even undermine the integrity of the asset being inspected or repaired – with potentially catastrophic consequences, e.g. accidental puncture of subsea pipelines. The robots themselves are high-value assets, whose loss or damage may be very costly. Moreover, there are increasing demands on regulating autonomous robots to build public trust in their use. Yet, the analysis of safety and reliability for autonomous robots poses a significant challenge due to the inevitable uncertainties in a mission. Potential sources of risk range from failures of sensors or hardware, to built-in algorithms making poor choices in a stochastic environment. Key industrial foresight reviews (Lane et al. 2016) outline a vision of *self-certifying*

robotic systems, i.e. systems that continuously monitor their current and predicted performance and assess it within a predefined certification framework (Robu, Flynn, and Lane 2018; Fisher et al. 2018).

Probabilistic model checking (PMC) (Kwiatkowska, Norman, and Parker 2018) has been successfully used to analyse quantitative properties of systems across a variety of application domains, including robotic systems (Luckcuck et al. 2018). This involves the construction of a probabilistic model, commonly using Discrete Time Markov Chain (DTMC), Continuous Time Markov Chain (CTMC) and Markov Decision Process (MDP) when considering non-deterministic actions, that formally represents the behaviour of a system over time. The properties of interest are normally specified in Linear Temporal Logic (LTL) or Probabilistic Computational Tree Logic (PCTL), then systematic exploration and analysis is performed to check if a claimed property holds.

One inherent problem for most (if not all) formal verification techniques is that the verification assumes the formal model (e.g. DTMC) accurately reflects the actual behaviour of the real-world system (Calinescu et al. 2016). Comparing to conventional systems - for which we might argue the formal model is fairly accurate – it becomes a tougher issue for systems in changing, unexpected environments and with autonomous features. To handle the issue, the runtime verification idea was proposed (Epifani et al. 2009; Calinescu et al. 2012) by keeping the formal model alive and continuously updated when seeing new data at runtime.

In this paper, we propose a tailored PMC framework for verifying safety-critical robots working in extreme environments. Firstly, we formalise how the robot works as a layered and parametric DTMC (transition probabilities are unknown parameters), then feed two Bayesian estimators for different types of parameters with operational data. For each mission, safety and reliability properties are verified at both:

**Pre-mission**, based on the best knowledge from previous similar missions, lab-simulations and experts, which provides assurance before the robot undertaking a mission.

**Runtime** (i.e. during a running mission), using a DTMC updated in real time. This allows the robot to take protection actions (e.g. restart with new control policies or abort the current mission) whenever a property of interest is being violated, which offers an additional control layer, independent

of the robot’s front-end planning engine. This follows the “defence in depth” design paradigm for safety critical systems, by providing an extra and diverse layer of protection.

To achieve these, we identify two types of transition parameters and introduce two novel Bayesian estimators as:

*A. Catastrophic failure related parameters*, which represent the probability of seeing a catastrophic failure when the robot is in an unsafe state. In practice, we cannot observe sufficient data of catastrophic failures to provide good estimations, since even if we do observe any, we normally will redesign/update the robot making the failure data obsolete. So effectively, for these model parameters, we will only collect *failure-free* data that lead to very optimistic estimations by existing Bayesian estimators. To avoid underestimating the chances of catastrophic failure, our claim is inference in such settings has to be carried out in a *conservative* way. The conservative Bayesian inference (CBI) method (Bishop et al. 2011; Strigini and Povyakalo 2013; Zhao et al. 2015; 2017) was developed for safety-critical software to answer the question what can be claimed rigorously about the reliability when seeing failure-free runs. To our knowledge, our work is the first to develop a CBI estimator for catastrophic failure parameters in robotics.

*B. Non-catastrophic failure related parameters*, which represent the transitions among normal, unsafe and non-catastrophic states, thus sufficient data can be collected as more missions are conducted. Bayesian methods yielding point estimates, e.g. (Epifani et al. 2009; Calinescu et al. 2014), are affected by unquantified estimation errors which will be propagated and compounded in the verification in ways that are unknown but likely to be significant, as highlighted in (Calinescu et al. 2016). Here we introduce an imprecise probability model with *sets of priors* (Walter and Augustin 2009; Walter, Aslett, and Coolen 2017) to (i) get upper and lower bounds on the posterior estimates whose range measures the estimation errors; (ii) allow modelling imperfect prior knowledge/data from experts/lab-simulations in a flexible way; and (iii) detect *prior-data conflict* (Evans and Moshonov 2006) when observe surprising data in a mission to provide protection from the robot’s epistemic limits.

We illustrate our new method with an example of an unmanned underwater vehicle (UUV) in the context of a valve turning scenario from the PANDORA<sup>2</sup> project which created UUVs that keep going under extreme uncertainty.

Next, we present the necessary background concepts. Two new Bayesian estimators are described in Section 3; Section 4 describes the new framework with an illustrative example. Section 5 summarises the related work. Contributions, limitations and future work are concluded in Section 6.

## 2 Preliminaries

### 2.1 DTMC and PCTL

Discrete Time Markov Chain (DTMC) is a widely-used model for formalising stochastic systems. From the verification point of view, DTMC can serve as a secondary protection model *given* an optimal policy (specifying a procedure for autonomous action selection). For instance, if the

primary planner synthesises an optimal policy via MDP and reinforcement learning (Pathak, Pulina, and Tacchella 2018; Pathak et al. 2013), then given the optimal policy, we obtain an *induced DTMC* (Puterman 2014) which can be revised to emphasise the safety and reliability aspects by, e.g. adding more transitions to states representing hazards.

**Definition 1.** A DTMC is a tuple  $(S, s_1, \mathbf{P}, L)$ , where:

- $S$  is a (finite) set of states; and  $s_1 \in S$  is an initial state;
- $\mathbf{P} : S \times S \rightarrow [0, 1]$  is a probabilistic transition matrix such that  $\sum_{s' \in S} \mathbf{P}(s, s') = 1$  for all  $s \in S$ ;
- $L : S \rightarrow 2^{AP}$  is a labelling function assigning to each state a set of atomic propositions from a set  $AP$ .

We use the notation  $p_{ij} = \mathbf{P}(s_i, s_j)$ , and  $i, j$  are integers in  $[1, k]$  by assuming there are  $k$  states in total. Given an optimal probabilistic<sup>3</sup> policy, the transition probability of the induced DTMC is defined as the total probability of:

$$p_{ij} = \sum_{a \in A} \pi_a(s_i) \cdot Pr(s_j | s_i, a) \quad (1)$$

where  $\pi_a(s_i)$  is the probability of executing action  $a$  in state  $s_i$  according to policy  $\pi$  and  $Pr(s_j | s_i, a)$  is the probability that  $s_j$  is the next state when action  $a$  is executed in state  $s_i$ . In this paper, we assume the optimal policy  $\pi$ , i.e.  $\pi_a(s_i)$ , is given by a separate planner and  $Pr(s_j | s_i, a)$  will be Bayesian updated via operational data.

The safety and reliability properties to be checked can be specified in Probabilistic Computation Tree Logic (PCTL).

**Definition 2.**  $AP$  is a set of atomic propositions and  $ap \in AP, p \in [0, 1], t \in \mathbb{N}$  and  $\bowtie \in \{<, \leq, >, \geq\}$ . The syntax of PCTL is defined by *state formulae*  $\Phi$  and *path formulae*  $\Psi$ .

$$\begin{aligned} \Phi &::= true \mid ap \mid \Phi \wedge \Phi \mid \neg \Phi \mid \mathcal{P}_{\bowtie p}(\Psi) \\ \Psi &::= X\Phi \mid \Phi U^{\leq t} \Phi \end{aligned}$$

where the temporal operator  $X$  is called *Next* and  $U$  is called *Until*. State formulae  $\Phi$  is evaluated to be either true or false in each state. Satisfaction relations for a state  $s$  are defined:

$$\begin{aligned} s &\models true \quad , \quad s \models ap \quad \text{iff} \quad ap \in L(s) \\ s &\models \neg \Phi \quad \text{iff} \quad s \not\models \Phi \\ s &\models \Phi_1 \wedge \Phi_2 \quad \text{iff} \quad s \models \Phi_1 \text{ and } s \models \Phi_2 \\ s &\models \mathcal{P}_{\bowtie p}(\Psi) \quad \text{iff} \quad Pr(s \models \Psi) \bowtie p \end{aligned}$$

$Pr(s \models \Psi) \bowtie p$  is the probability of the set of paths starting in  $s$  and satisfying  $\Psi$ . Given a path  $\psi$ , if denote its  $i$ -th state as  $\psi[i]$  and  $\psi[0]$  is the initial state. Then the satisfaction relation for a path formula for a path  $\psi$  is defined as:

$$\begin{aligned} \psi &\models X\Phi \quad \text{iff} \quad \psi[1] \models \Phi \\ \psi &\models \Phi_1 U^{\leq t} \Phi_2 \quad \text{iff} \quad \exists 0 \leq j \leq t \\ &\quad (\psi[j] \models \Phi_2 \wedge (\forall 0 \leq k < j \psi[k] \models \Phi_1)) \end{aligned}$$

It is worth mentioning that both DTMC and PCTL can be augmented with rewards/costs, cf. (Filieri and Tamburrelli 2013), which can be used to model, e.g. the energy/time

<sup>2</sup><http://persistentautonomy.com/?p=1436>

<sup>3</sup>Deterministic policy is a special case with probability 1.

consumption of the robot in a mission. Our approach is also compatible for those properties which we omit in this paper.

By formalising the robot and its required properties in DTMC and PCTL respectively, the verification focus shifts to the checking of *reachability* in a DTMC. In other words, PCTL expresses the constraints that must be satisfied concerning the probability, starting from the initial state, of reaching some states being labelled as, e.g. unsafe, failure or success. We use the tool PRISM (Kwiatkowska, Norman, and Parker 2011) which employs symbolic model checking algorithm to calculate the actual probability that a path formulae is satisfied (by extending the PCTL definition with  $\mathcal{P}_{=?}(\Psi)$ ), then comparing with a required bound if given.

## 2.2 Parametric model checking

PMC based on DTMC normally assumes the transition probabilities are known as constants which can be estimated from existing data and experts at design-time, or through system runtime monitoring. However, this traditional technique may not be suitable for runtime analysis in terms of the excessive time and power consumption, especially for robots. The idea of *parametric* model checking (ParaMC), proposed by (Daws 2005), provides an efficient solution (Filieri, Ghezzi, and Tamburrelli 2011). As shown by (Jansen et al. 2014), modern tools can solve parametric DTMCs with thousands of states and transitions using reasonable computing resources, beyond the required capability of our methods here.

ParaMC can analyse DTMC whose transition probabilities are specified as functions over a set of parameters, e.g. Equation (1). The result is then given as a closed-form rational function of the parameters which brings two advantages: (i) The verification can be divided into two steps. The computationally expensive symbolic analysis can be done at pre-mission stage when time and power constraints are not strong; then, during the mission, only substitutions are needed to replace the parameters in the closed-form expression with actual values learnt at runtime; (ii) Monotonicity analysis of the parameters can be easily done via the closed-form expression. Since, instead of point estimates, our new Bayesian estimator provides bounds for the parameters, thus monotonicity analysis is necessary to obtain meaningful bounds for the verification results.

## 3 Estimates for DTMC transition parameters

### 3.1 A fundamental Bayesian estimator

In a DTMC, given a current state  $i$ , the transition to a next state follows a *categorical distribution*. Due to the Markov property, i.e. the choice of a next state only depends on the current one, the categorical distributions for each state are *independent*. Hence, as we observe repeated transitions from the state  $i$ , the repeated categorical process follows a *multinomial distribution*. Now the problem is reduced to the *localised* learning of  $k$  independent multinomial distribution, where  $k$  is the number of states in the DTMC. From a Bayesian inference perspective, the posterior estimation requires a statistical model (the likelihood function) and a prior distribution. Note a complete description of Bayesian inference is beyond the scope of this paper.

For the  $i$ -th row of  $\mathbf{P}$ , if we observe the data of transition number from state  $i$  to  $j$  as  $n_{ij}$ , and  $n_i = \sum_{j=1}^k n_{ij}$  is the total number of outgoing transitions from state  $i$ , then the likelihood function is (by omitting the combinatorial factor which will be cancelled in the Bayesian formula):

$$Pr(\text{data} \mid p_{i1}, \dots, p_{ik}) = \prod_{j=1}^k p_{ij}^{n_{ij}} \quad (2)$$

The method in (Epifani et al. 2009) uses a Dirichlet distribution as priors for a given  $i$ -th row of  $\mathbf{P}$ :

$$(p_{i1}, \dots, p_{ik}) \sim Dir(n_i^{(0)} p_{i1}^{(0)}, \dots, n_i^{(0)} p_{ik}^{(0)}) \quad (3)$$

where  $n_i^{(0)} p_{i1}^{(0)}, \dots, n_i^{(0)} p_{ik}^{(0)}$  are the *canonical parameters* of the Dirichlet.  $p_{ij}^{(0)}$  is the prior expectation for the transition probability  $p_{ij}$ , and larger  $n_i^{(0)}$  leads to more concentrated probability measure around  $p_{ij}^{(0)}$ . Thus,  $n_i^{(0)}$  are quantifying the strength of beliefs in the prior  $p_{ij}^{(0)}$ , or a ‘‘pseudo-count’’ which can be interpreted as the size of an imaginary sample that gives the prior estimation (Walter and Augustin 2009).

Then applying the Bayes rule, and thanks to both the conjugacy (with the multinomial likelihood function) and the canonical form, the posterior with updated parameters are:

$$n_i^{(n_i)} = n_i^{(0)} + n_i \quad (4)$$

$$p_{ij}^{(n_i)} = \frac{n_i^{(0)}}{n_i^{(0)} + n_i} \cdot p_{ij}^{(0)} + \frac{n_i}{n_i^{(0)} + n_i} \cdot \frac{n_{ij}}{n_i} \quad (5)$$

Note, the upper index  $^{(0)}$  is used to identify the prior parameters, in contrast, the  $^{(n_i)}$  denotes the posterior parameters after observing  $n_i$  outgoing transitions from the state  $i$ . As Equation (5) shows, after seeing  $n_{ij}$  out of  $n_i$  as data, the posterior  $p_{ij}^{(n_i)}$  is a *weighted* sum of two terms: the prior estimate of  $p_{ij}^{(0)}$  and the  $n_{ij}/n_i$  which is the frequency of the relevant transition in the data. The weights are proportional to the  $n_i^{(0)}$  (the ‘‘pseudo-count’’ of prior simple size) and  $n_i$  (the ‘‘actual-count’’ of data sample size). Smaller  $n_i^{(0)}$  represents lower confidence in the priors and the runtime data will dominate the posteriors. When  $n_i^{(0)} \simeq 0$ , (5) reduces to the Maximum Likelihood Estimation (Epifani et al. 2009).

### 3.2 A Bayesian estimator using sets of priors

There are at least two practical issues with the method in Section 3.1: (i) how to justify a particular choice of prior model parameters, i.e.  $n_i^{(0)}$  and  $p_{ij}^{(0)}$ . In other words, whether one can truly express the subjective and imprecise prior knowledge with the *exactness* that a particular prior distribution requires; and (ii) how to measure the error of single point estimation which will be propagated and compounded in the later model checking in ways that are unknown but likely to be significant, as pointed in (Calinescu et al. 2016).

To address these two concerns, we utilise an imprecise probability model using sets of priors (Walter and Augustin 2009; Walter, Aslett, and Coolen 2017) to model more vague

prior knowledge by eliciting bounds of the prior parameters, and also resulting bounds for the posteriors whose range measures the estimation errors. Also importantly, it is sensitive to detect the *prior-data conflict*, i.e. conflicts between prior assumptions and observed data (Evans and Moshonov 2006), which is useful to alter dangerous situations and will be discussed later.

To be exact, instead of a single value, we elicit an interval for each prior parameters in (3) and denote these as:

$$n_i^{(0)} \in \left[ \underline{n}_i^{(0)}, \overline{n}_i^{(0)} \right] \quad , \quad p_{ij}^{(0)} \in \left[ \underline{p}_{ij}^{(0)}, \overline{p}_{ij}^{(0)} \right]$$

Then as proved in (Walter and Augustin 2009), the lower and upper bounds for the posteriors of interest  $p_{ij}^{(n_i)}$  are:

$$\underline{p}_{ij}^{(n_i)} = \begin{cases} \frac{\overline{n}_i^{(0)} p_{ij}^{(0)} + n_{ij}}{\overline{n}_i^{(0)} + n_i} & \text{if } \frac{n_{ij}}{n_i} \geq \underline{p}_{ij}^{(0)} \\ \frac{\underline{n}_i^{(0)} p_{ij}^{(0)} + n_{ij}}{\underline{n}_i^{(0)} + n_i} & \text{if } \frac{n_{ij}}{n_i} < \underline{p}_{ij}^{(0)} \end{cases} \quad (6)$$

$$\overline{p}_{ij}^{(n_i)} = \begin{cases} \frac{\overline{n}_i^{(0)} \overline{p}_{ij}^{(0)} + n_{ij}}{\overline{n}_i^{(0)} + n_i} & \text{if } \frac{n_{ij}}{n_i} \leq \overline{p}_{ij}^{(0)} \\ \frac{\underline{n}_i^{(0)} \overline{p}_{ij}^{(0)} + n_{ij}}{\underline{n}_i^{(0)} + n_i} & \text{if } \frac{n_{ij}}{n_i} > \overline{p}_{ij}^{(0)} \end{cases} \quad (7)$$

When  $\frac{n_{ij}}{n_i} \notin \left[ \underline{p}_{ij}^{(0)}, \overline{p}_{ij}^{(0)} \right]$ , i.e. the prior-data conflict is at hand, the posterior interval  $\left[ \underline{p}_{ij}^{(n_i)}, \overline{p}_{ij}^{(n_i)} \right]$  becomes wider, meaning we are *even less certain* about the posterior estimation comparing to the priors. This is a new property comparing to other imprecise probability models in which the range of the posterior interval will *always*, regardless of the prior-data conflict, decrease (i.e. converge to the observed data) as the sample size increases.

### 3.3 Conservative Bayesian inference

Catastrophic failures are modelled in our new framework by imposing transitions from unsafe states to catastrophic failures of different modes, e.g. the parameter  $x$  in Fig.1 as an instance. The “true unknown” values of these transition parameters may lie in very small orders of magnitude, say  $10^{-5} \rightarrow 10^{-9}$ , as a result of rigorous development process and safety-critical designs. This poses a great challenge for the estimation of such smaller failure rates in terms of that infeasible amount of statistical testing or operational time is required to observe sufficient failure data (Littlewood and Strigini 1993). More practically, even if we did observe any catastrophic failures, we normally will redesign/update the robots, which makes the failure data obsolete. So effectively, we will only collect *catastrophic-failure-free* data. Such “good news” would increase our confidence that the robot has a smaller chance to cause catastrophic failures, hence our claim is such inference has to be done in a *conservative* way. The conservative Bayesian inference (CBI) method (Bishop et al. 2011; Strigini and Povyakalo 2013; Zhao et al. 2015; 2017) was developed for safety-critical software to answer what can be claimed rigorously about the reliability when seeing failure-free runs. Here we introduce CBI as our Bayesian estimator for the catastrophic failure related parameters. Note, we only present the more practical case of seeing no catastrophic failures in this paper. The essential CBI can be extended to model very scarce failures.

In a given unsafe state  $i$ , as discussed above, the outgoing transitions follow a multinomial distribution, so *marginally*, the number of transitions to the catastrophic failure state, say  $n_{ij}$ , is a binomial one,  $n_{ij} \sim Bin(n_i, x)$  where  $x$  is the transition probability and  $n_i$  is the total number of outgoing transitions from state  $i$ . Then the likelihood function for the data of no catastrophic failure observed, i.e.  $n_{ij} = 0$ , is:  $Pr(\text{data}|x) = (1 - x)^{n_i}$ . So if we have a prior distribution for  $x$ , say  $f(x)$ , by Bayes rule we know the posterior expectation is:

$$E(x | \text{data}) = \frac{\int_0^1 x(1-x)^{n_i} f(x) dx}{\int_0^1 (1-x)^{n_i} f(x) dx} \quad (8)$$

Conventionally, starting from Equation (8), we assign a parametric family for  $f(x)$ , e.g. a conjugate Beta in this case, or like the Dirichlet-Multinomial case in Section 3.1. However, the use of conjugacy is based on the assumption that the practical situations we are dealing with have large quantities of failure data, in which the dominant contribution to posterior belief via Bayes Theorem comes from the likelihood function, i.e. situations in which “the data can speak for themselves”. We do not have this luxury for safety-critical systems with very limited or no failure data, so any use of a particular parametric family for  $f(x)$  is questionable.

Instead of assuming a *complete* prior distribution that follows a parametric family, the assessors are more likely to have (or be able to justify) some very limited *partial* prior knowledge, e.g. two possible scenarios: (i) “I am 80% confident the robot will not have any catastrophic failure in this unsafe state” is a confidence in its catastrophic-failure-freeness, i.e.  $Pr(x = 0) = 0.8$ . See (Littlewood and Rushby 2012; Strigini and Povyakalo 2013) for the arguments of such partial prior knowledge; (ii) “I am 90% confident the probability of seeing a catastrophic failure from this unsafe state is smaller than 0.001” is a confidence bound on a given probability of seeing catastrophic failure, i.e.  $Pr(x \leq 0.001) = 0.9$ . Such partial prior knowledge could be supported by, e.g. when evidence is presented showing the system is strictly developed against IEC61508 SIL3<sup>4</sup>.

In the above scenarios, the elicited partial prior knowledge is far from a complete prior distribution. Rather, if treat the partial priors as *constraints* on a distribution, then there must be an *infinite set of prior distributions* satisfying the prior constraints. Note, this set of priors is different from the one in Section 3.2 which still assumes a parametric family (i.e. Dirichlet). Now the problem reduces to find the *most conservative* prior distribution (in the sense of giving a maximum posterior expected transition probability) in the infinite set of priors satisfying the elicited prior constraints.

For example, as the first scenario, the assessor only has a  $\theta$  confidence in its catastrophic-failure-freeness:

$$Pr(x = 0) = \theta \quad (9)$$

As proved in (Strigini and Povyakalo 2013), to maximise (8), the corresponding  $f(x)$ , that subjects to the constraint (9), is a two-point one with probability mass at  $Pr(x =$

<sup>4</sup>[http://en.wikipedia.org/wiki/Safety\\_integrity\\_level](http://en.wikipedia.org/wiki/Safety_integrity_level).

0) =  $\theta$  and  $Pr(x = q) = 1 - \theta$ , where  $q$  is an optimisation point that can be obtained numerically. Thus, by such two-point  $f(x)$ , Equation (8) can be further bounded by:

$$\begin{aligned} E(x | \text{data}) &\leq \frac{(1 - \theta)q(1 - q)^{n_i}}{\theta + (1 - \theta)(1 - z)^{n_i}} < \frac{(1 - \theta)q(1 - q)^{n_i}}{\theta} \\ &\leq \frac{(1 - \theta)}{\theta(n_i + 1)} \left(1 - \frac{1}{n_i + 1}\right)^{n_i} \end{aligned} \quad (10)$$

Note the last step above is because that  $q(1 - q)^{n_i}$  reaches its maximum when  $q = 1/(n_i + 1)$ .

Depending on the particular form of partial prior knowledge elicited, the worst-case priors varies, so does the posteriors. For cases of other forms of partial prior knowledge, see (Bishop et al. 2011; Zhao et al. 2015; 2017).

To summarise, in this section, we present two advanced Bayesian estimators for the two types of transition parameters. Mixed use of incorrect estimators will lead to either “too conservative to be useful results” or “continuous prior-data conflict”, which we omit here due to the page limit.

## 4 The framework with an UUV example

### 4.1 The running example

We use an underwater search and valve turning mission in which a UUV equipped with electrical manipulators, stereo cameras, a specifically designed end-effector which had a camera in-hand with force and torque sensors to (i) locate a valve panel among different locations; and (ii) modify the valve handles to achieve different panel configurations. Possible disturbances are, e.g. muddy water, strong currents, blocked valves and panel occlusion. We formalise a single valve operation as an instance of the DTMC in Fig. 1 which has been simplified with only 6 states, whilst more realistic formal models, e.g. with more states and transitions as shown as the dotted shapes in the figure, can be obtained in our proposed generic way as follows.

The DTMC in Fig. 1 has 3 layers. The normal operation and safe states are modelled in the first layer, in which each state can transit into one or more unsafe states with some probabilities, e.g. in the  $S1$  state of searching valve panel, the UUV may encounter a region of muddy water leading to the unsafe low visibility state  $S4$ . We group all the unsafe states in the middle layer which can either transit back to the normal operations or cause various modes of failures, including catastrophic ones, in the bottom layer, e.g. the propeller has a higher risk of malfunction when the UUV is in a muddy water. We believe the structure of a realistic DTMC can be built in a generic way that the top layer can be derived from the original policy-making model, e.g. MDP and reinforce learning as in (Pathak, Pulina, and Tacchella 2018; Pathak et al. 2013), and the bottom two layers can be obtained from traditional safety and reliability analysis based on expert knowledge like Failure Mode, Mechanism and Effect Analysis (FMMEA) and Fault Tree Analysis (FTA).

Practically, most of the transition probabilities in the formal Markov model can be argued as known constants, e.g. by prognostics on the remaining useful life of key components (Romain et al. 2017), leaving only a few transition

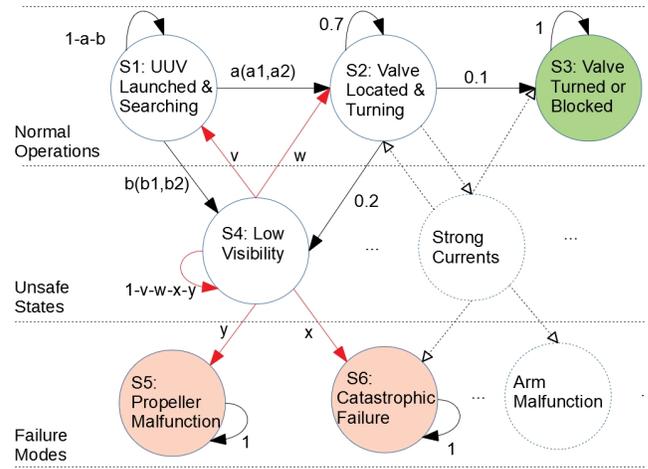


Figure 1: A layered DTMC modelling an UUV in a valve turning mission. The dotted shapes illustrate possible extensions to make the simplified example more realistic.

probabilities as unknown parameters, e.g. the  $x$  and  $y$  in Fig. 1 (thus a parametric DTMC), which will be Bayesian estimated at runtime with collected data.

Fig. 1 is also an *induced* DTMC, since we assume the *given* optimal policy is a probabilistic one for  $S1$ , e.g. with  $\gamma$  probability to do speed-1 and  $(1 - \gamma)$  probability to do speed-2, where  $\gamma$  is learnt by a separate primary planner. Then the induced transition probability from  $S1$  to  $S2$  is a function of  $a_1, a_2$  as given by Equation (1):  $a(a_1, a_2) = \gamma a_1 + (1 - \gamma) a_2$  where  $a_1$  and  $a_2$  are the probabilities of transiting from  $S1$  to  $S2$  with autonomous speed-1 and speed-2 respectively. Similarly for  $b(b_1, b_2)$ . For  $S4$ , we assume a deterministic policy (e.g. always do safe mode speed in low visibility), thus the outgoing transition probabilities are single parameters.

We are interested in 4 requirements: **R1**: What is the probability of completing a *next* mission? **R2**: What is the probability of seeing catastrophic failures in a *next* mission? **R3**: What is the probability of completing *this* mission? **R4**: What is the probability of seeing catastrophic failures in *this* mission? The R1 and R2 will be analysed at the pre-mission stage, whose PCTL with the initial state specified as  $S1$  are:

$$P_{=?} [ \text{true } U(s = 3) ] \quad , \quad P_{=?} [ \text{true } U(s = 6) ]$$

Whilst, R3 and R4 will be verified at runtime. Their PCTL are same as above but with different specified initial states.

### 4.2 The framework and simulated experiments

As shown in Fig. 2, our framework has two stages:

At the *pre-mission stage*, a separate planner will first synthesise the optimal policy, which is assumed as given in the current framework. Given the optimal policy, an induced DTMC is derived and then revised to emphasise the safety and reliability by adding transitions to states representing hazards and various modes of failures, in which some transition probabilities are unknown parameters. Then ParaMC is conducted by tools like PRISM to generate rational functions for each PCTL property. The ParaMC results will be

recorded and also reused by the on-board verifier to speed up the verification at runtime. Monotonicity analysis for each parameter in the rational functions is necessary due to the use of bounds for each transition parameter. Data analysis of experts’ knowledge, lab-simulations and previous similar missions is performed to form prior knowledge required by the two Bayesian estimators. Finally, by substituting parameters in the rational functions with the prior estimates, pre-mission verification results are obtained, which represents our best efforts of assurance before launching the robot.

In mission (runtime), a monitor collects new data and does Bayesian estimates on the transition parameters. The on-board verifier checks if any required properties are being violated at runtime based on the ParaMC results at pre-mission stage. When any violation presents, actions like abort or instantly restart the mission (with a repaired policy) can be taken, whose implementation is beyond the scope of this paper. Otherwise, the UUV will continue and keep monitoring and verifying until the mission is complete.

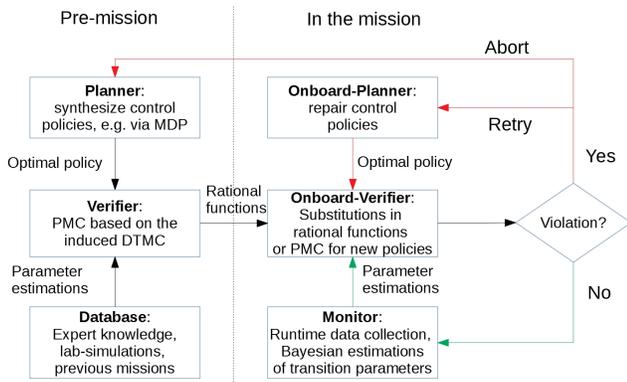


Figure 2: Overview of the verification framework.

We demonstrate and evaluate our framework by simulated experiments based on the UUV example of Fig. 1. First, to simulate and collect data of both previous missions and the current one for the pre-mission and runtime verifications respectively, we use the PRISM simulation module by assuming the *true unknown* transitions probabilities in Fig. 1 are:  $a_1 = 0.05, a_2 = 0.03, b_1 = 0.2, b_2 = 0.1, x = 10^{-5}, y = 0.001, v = 0.3, w = 0.01$ . For the optimal policy of state 1, we assign a random number in  $[\frac{25}{40}, \frac{35}{40}]$ <sup>5</sup> for  $\gamma$  in each simulated mission to mimic the practical case of changing policies. Then, it results in 49 previous missions and a 50-th as the current one<sup>6</sup> which explicitly using  $\gamma = 0.75$  as the optimal probabilistic policy of state 1.

At the pre-mission stage, given the data of previous 49 missions (e.g. 1566 outgoing transitions from  $S_1$  with the action of speed-1, in which 1180 loops in  $S_1$ , 85 to  $S_2$  and 301 to  $S_4$ ), together with the prior knowledge in Table 1, the posteriors by our Bayesian estimators are also listed there.

<sup>5</sup>All assumed values of the parameters in the simulation refer to the information provided by the PANDORA project.

<sup>6</sup>Without loss of generality, we choose a 50-th mission with relatively more transitions (522 to be exact) for a better illustration.

Next, given  $\gamma = 0.75$  and using the ParaMC engine of PRISM, we obtain closed-form results for R1 and R2 (also R3 and R4 with different initial states for the use at runtime) as rational functions whose monotonicity can be easily analysed with respect to the 8 transition parameters. Then we obtain the pre-mission verification results:

$$R1 \in [0.857, 0.961], \quad R2 \in [7.96 \times 10^{-4}, 1.55 \times 10^{-3}]$$

In the current 50-th mission, the posteriors of the 8 transition parameters in Table 1 are in turn used as priors for the runtime Bayesian updates. The real-time estimates at each discrete time step of the transitions are plotted in (a)-(h) of Fig. 3, correspondingly R3 and R4 are plotted in (i) and (j).

As shown in the (a) of Fig. 3, the CBI estimation for  $x$  decreases along with the execution of the robot without catastrophic failures. For non-catastrophic failure related parameters of (b)-(h), instead of using CBI to obtain a conservative point estimation, we plot both upper and lower bounds via Bayesian inference with sets of priors. Basically for each of these parameters, we first observe a divergence trend of the two bounds which then start to converge as more data is collected. Indeed, at the beginning of a mission, due to the sparse data collected, the prior-data conflict phenomenon is normal and foreseeable (e.g. 3 heads out of 3 tosses of a fair coin is conflicting our prior belief in its fairness). Such conflict is reflected by a wider bounded (i.e. less certain) estimation for each parameter and then propagated to the overall verification results.

For instance, verifications on both R3 and R4 become less and less certain in the first 100 steps, partly due to the prior-data conflict of  $a_1$ . Then, the 101-th step is a transition from  $S_1$  to  $S_2$  with speed-1 action, which not only provides a new estimate of  $a_1$  but also resolves the prior-data conflict to some extent, leading to tighter bounds for  $a_1, R3$  and R4.

As more runtime data is observed, the general trend for each pair of bounds in Fig. 3 starts to converge, meaning a more certain verification result. Note, both R3 and R4 highly depend on which state the UUV is currently in, so the obvious “bumps” in the plot for R3 are because the UUV is in  $S_2$  which certainly has a relatively higher chance to complete the mission. Similarly for the bumps in the plot of R4.

The above example is an ideal case in the sense that the new learning agrees with what it learnt before, i.e. all the 50 missions are simulated based on an MDP with the same transition probabilities, like the probability of transiting back to the normal state  $S_1$  from the low visibility state  $S_4$  is constantly assumed as  $v = 0.3$  for all 50 missions. For exceptional cases, e.g. there is a large body of muddy water implying  $v = 0.8$ , our method will detect them as prior-data conflict through the whole mission, thus a continuous divergence of the bounds of related transition parameters will be observed, so consequently does an overall divergence of the verification results R3 and R4. We label this exceptional case as “known unknowns” in the sense that the new unknown (e.g.  $v = 0.8$ ) is learnt in an *informed* way such that we know how much contradict it is to our prior knowledge, thus leading a less certain (i.e. wider bounds) new estimation.

We also would like to highlight the exceptional case of “unknown unknowns” in which the robot fails to know

Table 1: Prior knowledge and posterior estimations for the transition parameters given the data of 49 previous missions.

	$x$	$y$	$v$	$w$	$a_1$	$b_1$	$a_2$	$b_2$
pse. cou.	N/A	[100, 300]	[100, 300]	[100, 300]	[100, 300]	[100, 300]	[50, 100]	[50, 100]
pri. est.	$\theta = 0.9$	[0.001, 0.01]	[0.1, 0.4]	[0.001, 0.01]	[0.01, 0.1]	[0.1, 0.5]	[0.01, 0.1]	[0.1, 0.5]
post. est.	$3.2e-5$	[0.0014, 0.0032]	[0.27, 0.33]	[0.0097, 0.012]	[0.047, 0.062]	[0.18, 0.24]	[0.03, 0.04]	[0.09, 0.16]

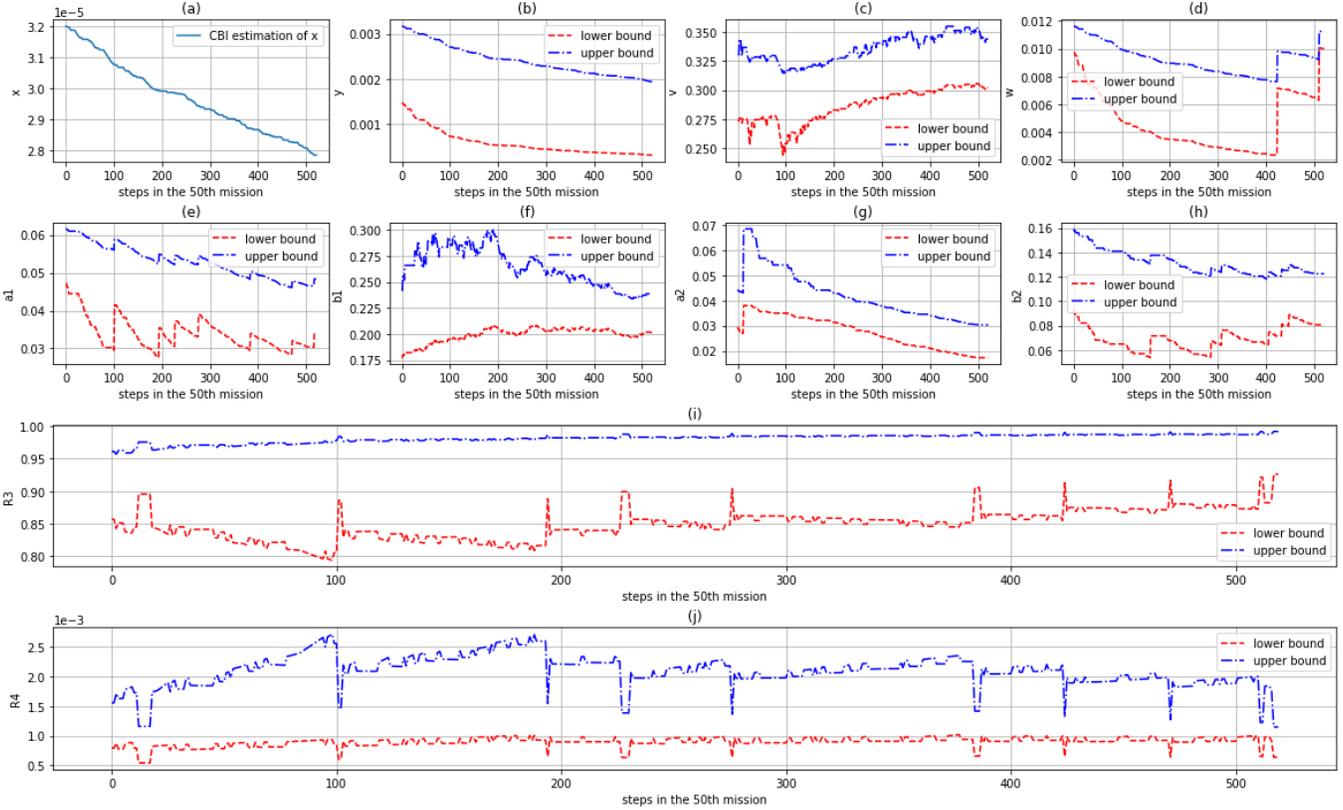


Figure 3: Runtime Bayesian estimations for the 8 transition parameters (a)-(h) and verification results for R3 (i) and R4 (j).

which state it is in and *without knowing this fact*. For example, the UUV now is in a muddy water (i.e.  $S_4$ ), but the sensor fails to detect this unsafe state so that the UUV believes it is still in the normal state  $S_1$ . As a consequence, the UUV does a wrong action of either speed-1 or speed-2 (i.e. the probabilistic policy in  $S_1$ ) in the unsafe state  $S_4$  instead of the safe speed (i.e. the deterministic policy for  $S_4$ ). Our method is also able to alert such dangerous case by detecting prior-data conflict phenomenon that happens *simultaneously* for many parameters. Because, in the example above, although the UUV is *actually* in  $S_4$ , it will *experience* constant loops in  $S_1$  without transiting to neither  $S_4$  (as the sensors fail to detect the abnormal states) nor  $S_2$  (as wrong actions are taken, assuming going fast in a muddy water will never detect the valves). Consequently, prior-data conflict happens simultaneously for all  $a_1, a_2, b_1, b_2$  parameters.

More examples for the exceptional cases of “known unknowns” and “unknown unknowns” will be discussed in future due to the page limits here. We believe, detecting the prior-data conflict effect during the mission can assure pro-

tection against from faulty knowledge (i.e. epistemic limits) of the robot about its own state of health and environments.

## 5 Related work

How should autonomous systems be verified is a new challenging question along with their increasing applications (Fisher, Dennis, and Webster 2013). Formal methods must be integrated in order to develop, verify and provide certification evidence for large-scale and complex autonomous systems like robots (Farrell, Luckcuck, and Fisher 2018).

Model checking is a widely used formal method in verifying robotic systems, due to its relative simplicity and powerful automatic tools (Luckcuck et al. 2018). For instance, in (Webster et al. 2014), a proof-of-concept approach is presented to generate certification evidence for autonomous unmanned aircraft based on both model checking and simulation. PMC, as a variant, emphasises the inherent uncertainties of the formalised system. In (Konur, Dixon, and Fisher 2012; Gainer, Dixon, and Hustadt 2016), the complex and uncertain behaviours of robot swarms are

analysed by PMC. In (Norman, Parker, and Zou 2017; Pathak, Pulina, and Tacchella 2018), PMC is used to verify the control policies of robots in partially unknown environments. In a hostile environment, the movements of adversaries are modelled probabilistically in (Cizelj et al. 2011). The reliability and performance of UUVs is guaranteed by PMC in (Gerasimou, Calinescu, and Banks 2014).

Although runtime PMC is effective for assuring the quality of service-based systems (Calinescu et al. 2011) and self-adaptive systems (Calinescu et al. 2012; Filieri and Tamburrelli 2013), there is little research on runtime PMC for robots. In the UUV domain, the first work of runtime PMC is credited to (Gerasimou, Calinescu, and Banks 2014). However, it focuses on improving the scalability of runtime PMC by using software engineering techniques, which is also applicable to our work here that focuses on developing new methods of learning model parameters.

One of the initial methods to learn the transition probabilities of DTMC is in (Epifani et al. 2009), which later has been retrofitted for CTMC (Filieri, Ghezzi, and Tamburrelli 2012) and extended with ageing factors of collected data to accurately estimate time-varying transition probabilities (Calinescu et al. 2014). To reduce the noise and provide smooth estimates, a lightweight adaptive filter is proposed in (Filieri, Grunske, and Leva 2015). Whilst, above mentioned approaches yield point estimations, these can be affected by unquantified and potentially significant errors. The work in (Calinescu et al. 2016) is the first to synthesise bounds for unknown transition parameters. However, it is based on the theory of simultaneous confidence intervals, which is fundamentally different to the Bayesian approach presented here which has the distinct advantage of being able to embed various forms of prior knowledge.

## 6 Conclusions & future work

In this paper, we present a new framework to utilise PMC to assess the safety and reliability of robots at both pre-mission stage and runtime. Our main contributions are:

1. CBI is introduced with new closed-form results as a novel estimator for catastrophic failure related parameters.
2. Imprecise probability with sets of priors is introduced as another novel estimator for transition parameters. It allows to not only quantify the estimation errors and flexibly model imperfect prior knowledge, but also detect prior-data conflict to alter various dangerous situations.
3. A generic way is discussed to structure Markov models into layers to emphasise the system safety and reliability.
4. A real-world application of UUVs has been formalised which can be extended and reused for future research.

For illustration purpose, we demonstrate our methods with a stylised example which can be easily extended in our proposed generic way without changing the main conclusions, e.g. by considering the geographical waypoints in the top layer of the DTMC or listing complete failure modes in the bottom layer. The practicality of our new approach needs to be further evaluated with more case studies. We see potential issues like (i) insufficient safety analysis in

FTA/FMMEA to generate sound DTMC in the bottom two layers; and (ii) unnecessarily complex DTMC model with too many unknown transition parameters which require too much prior knowledge from experts and too much data to be collected at runtime that burdens the on-board sensors.

We plan to exploit more use of prior-data conflict and also *strong prior-data agreement* (Walter and Coolen 2016) to reflect the robot's epistemic limitations. Requirements to be verified for a robot should come from a higher level, e.g. when a robot is part of the Prognostics and Health Management (PHM) system, a PHM level PMC can be done in future to answer what to certify a robot. We also plan to propose a *lightweight* on-board re-planner based on an MDP with the up-to-date bounded transition parameters.

## 7 Acknowledgements

This work was supported by the UK EPSRC, through the Offshore Robotics for Certification of Assets (ORCA) Hub [EP/R026173/1]. We thank Lorenzo Strigini, Peter Bishop and Andrey Povyakalo from City, University of London who provided insights on the initial ideas of the work.

## References

- Bishop, P.; Bloomfield, R.; Littlewood, B.; Povyakalo, A.; and Wright, D. 2011. Toward a formalism for conservative claims about the dependability of software-based systems. *IEEE Tran. on Software Engineering* 37(5):708–717.
- Calinescu, R.; Grunske, L.; Kwiatkowska, M.; Mirandola, R.; and Tamburrelli, G. 2011. Dynamic QoS management and optimization in service-based systems. *IEEE Tran. on Software Engineering* 37(3):387–409.
- Calinescu, R.; Ghezzi, C.; Kwiatkowska, M.; and Mirandola, R. 2012. Self-adaptive software needs quantitative verification at runtime. *Comm. of the ACM* 55(9):69–77.
- Calinescu, R.; Rafiq, Y.; Johnson, K.; and Bakır, M. E. 2014. Adaptive model learning for continual verification of non-functional properties. In *Proc. of the 5th Int. Conf. on Performance Engineering*, 87–98. NY, USA: ACM.
- Calinescu, R.; Ghezzi, C.; Johnson, K.; Pezzé, M.; Rafiq, Y.; and Tamburrelli, G. 2016. Formal verification with confidence intervals to establish quality of service properties of software systems. *IEEE Tran. on Reliability* 65(1):107–125.
- Cizelj, I.; Ding, X.; Lahijanian, M.; Pinto, A.; and Belta, C. 2011. Probabilistically safe vehicle control in a hostile environment. *IFAC Proc. Volumes* 44(1):11803 – 11808.
- Daws, C. 2005. Symbolic and parametric model checking of discrete-time markov chains. In *Theoretical Aspects of Computing*, 280–294. Berlin, Heidelberg: Springer.
- Epifani, I.; Ghezzi, C.; Mirandola, R.; and Tamburrelli, G. 2009. Model evolution by run-time parameter adaptation. In *Proc. of the 31st Int. Conf. on Software Engineering*, 111–121. Washington, DC, USA: IEEE.
- Evans, M., and Moshonov, H. 2006. Checking for prior-data conflict. *Bayesian Analysis* 1(4):893–914.
- Farrell, M.; Luckcuck, M.; and Fisher, M. 2018. Robotics and integrated formal methods: Necessity meets opportu-

- nity. In *Proc. of the 14th Int. Conf. on Integrated Formal Methods*, 161–171. Cham: Springer.
- Filieri, A., and Tamburrelli, G. 2013. Probabilistic verification at runtime for self-adaptive systems. In *Assurances for Self-Adaptive Systems: Principles, Models, and Techniques*. Berlin, Heidelberg: Springer Berlin Heidelberg. 30–59.
- Filieri, A.; Ghezzi, C.; and Tamburrelli, G. 2011. Run-time efficient probabilistic model checking. In *Proc. of the 33rd Int. Conf. on Software Eng.*, 341–350. NY, USA: ACM.
- Filieri, A.; Ghezzi, C.; and Tamburrelli, G. 2012. A formal approach to adaptive software: continuous assurance of non-functional requirements. *Formal Aspects of Computing* 24(2):163–186.
- Filieri, A.; Grunske, L.; and Leva, A. 2015. Lightweight adaptive filtering for efficient learning and updating of probabilistic models. In *Proc. of the 37th Int. Conf. on Software Engineering*, 200–211. Piscataway, NJ, USA: IEEE Press.
- Fisher, M.; Collins, E.; Dennis, L.; Luckcuck, M.; Webster, M.; Jump, M.; Page, V.; Patchett, C.; Dinmohammadi, F.; Flynn, D.; Robu, V.; and Zhao, X. 2018. Verifiable Self-Certifying Autonomous Systems. In *2018 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*, 341–348.
- Fisher, M.; Dennis, L.; and Webster, M. 2013. Verifying autonomous systems. *Comm. of the ACM* 56(9):84–93.
- Gainer, P.; Dixon, C.; and Hustadt, U. 2016. Probabilistic model checking of ant-based positionless swarming. In *Towards Autonomous Robotic Systems*, 127–138. Cham: Springer International Publishing.
- Gerasimou, S.; Calinescu, R.; and Banks, A. 2014. Efficient runtime quantitative verification using caching, lookahead, and nearly-optimal reconfiguration. In *Proc. of the 9th Int. Symp. on Software Engineering for Adaptive and Self-Managing Systems*, 115–124. NY, USA: ACM.
- Jansen, N.; Corzilius, F.; Volk, M.; Wimmer, R.; Ábrahám, E.; Katoen, J.-P.; and Becker, B. 2014. Accelerating parametric probabilistic verification. In *Quantitative Evaluation of Systems*, 404–420. Springer International Publishing.
- Konur, S.; Dixon, C.; and Fisher, M. 2012. Analysing robot swarm behaviour via probabilistic model checking. *Robotics and Autonomous Systems* 60(2):199 – 213.
- Kwiatkowska, M.; Norman, G.; and Parker, D. 2011. PRISM 4.0: Verification of probabilistic real-time systems. In *Proc. of the 23rd Int. Conf. on Computer Aided Verification*, volume 6806 of *LNCS*, 585–591. Springer.
- Kwiatkowska, M.; Norman, G.; and Parker, D. 2018. Probabilistic model checking: Advances and applications. In *Formal System Verification: State-of-the-Art and Future Trends*. Cham: Springer International Publishing. 73–121.
- Lane, D.; Bisset, D.; Buckingham, R.; Pegman, G.; and Prescott, T. 2016. New foresight review on robotics and autonomous systems. Technical Report No. 2016.1, Lloyd’s Register Foundation, London, U.K.
- Littlewood, B., and Rushby, J. 2012. Reasoning about the reliability of diverse two-channel systems in which one channel is “possibly perfect”. *IEEE Tran. on Software Engineering* 38(5):1178–1194.
- Littlewood, B., and Strigini, L. 1993. Validation of ultra-high dependability for software-based systems. *Communications of the ACM* 36:69–80.
- Luckcuck, M.; Farrell, M.; Dennis, L.; Dixon, C.; and Fisher, M. 2018. Formal specification and verification of autonomous robotic systems: A survey. *arXiv preprint arXiv:1807.00048*.
- Norman, G.; Parker, D.; and Zou, X. 2017. Verification and control of partially observable probabilistic systems. *Real-Time Systems* 53(3):354–402.
- Pathak, S.; Pulina, L.; Metta, G.; and Tacchella, A. 2013. Ensuring safety of policies learned by reinforcement: Reaching objects in the presence of obstacles with the iCub. In *2013 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 170–175.
- Pathak, S.; Pulina, L.; and Tacchella, A. 2018. Verification and repair of control policies for safe reinforcement learning. *Applied Intelligence* 48(4):886–908.
- Puterman, M. L. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Robu, V.; Flynn, D.; and Lane, D. 2018. Train robots to self-certify their safe operation. *Nature* 553(7688):281.
- Romain, D.; Dickie, R.; Robu, V.; and Flynn, D. 2017. A review of the role of prognostics in predicting the remaining useful life of assets. In *Safety and Reliability - Theory and Applications*, 135. CRC Press.
- Strigini, L., and Povyakalo, A. A. 2013. Software fault-freeness and reliability predictions. In *Proc. of the 32nd International Conf. on Computer Safety, Reliability and Security*, volume 8153, 106–117. Springer.
- Walter, G., and Augustin, T. 2009. Imprecision and prior-data conflict in generalized bayesian inference. *Journal of Statistical Theory and Practice* 3(1):255–271.
- Walter, G., and Coolen, F. P. A. 2016. Sets of priors reflecting prior-data conflict and agreement. In *Information Processing and Management of Uncertainty in Knowledge-Based Systems*, 153–164. Cham: Springer.
- Walter, G.; Aslett, L.; and Coolen, F. P. A. 2017. Bayesian nonparametric system reliability using sets of priors. *International Journal of Approximate Reasoning* 80:67 – 88.
- Webster, M.; Cameron, N.; Fisher, M.; and Jump, M. 2014. Generating certification evidence for autonomous unmanned aircraft using model checking and simulation. *Journal of Aerospace Information Systems* 11(5):258–279.
- Zhao, X.; Littlewood, B.; Povyakalo, A.; and Wright, D. 2015. Conservative claims about the probability of perfection of software-based systems. In *Proc. of the 26th Int. Symp. on Software Reliability Engineering*, 130–140. IEEE.
- Zhao, X.; Littlewood, B.; Povyakalo, A.; Strigini, L.; and Wright, D. 2017. Modeling the probability of failure on demand (pfd) of a 1-out-of-2 system in which one channel is “quasi-perfect”. *Reliability Engineering & System Safety* 158:230–245.