communications engineering

ARTICLE

https://doi.org/10.1038/s44172-024-00162-y

OPEN

# Bayesian learning for the robust verification of autonomous robots

Xingyu Zhao[1✉], Simos Gerasimou[2✉], Radu Calinescu[2], Calum Imrie[2], Valentin Robu[3,4] & David Flynn[5]

Autonomous robots used in infrastructure inspection, space exploration and other critical missions operate in highly dynamic environments. As such, they must continually verify their ability to complete the tasks associated with these missions safely and effectively. Here we present a Bayesian learning framework that enables this runtime verification of autonomous robots. The framework uses prior knowledge and observations of the verified robot to learn expected ranges for the occurrence rates of regular and singular (e.g., catastrophic failure) events. Interval continuous-time Markov models defined using these ranges are then analysed to obtain expected intervals of variation for system properties such as mission duration and success probability. We apply the framework to an autonomous robotic mission for underwater infrastructure inspection and repair. The formal proofs and experiments presented in the paper show that our framework produces results that reflect the uncertainty intrinsic to many real-world systems, enabling the robust verification of their quantitative properties under parametric uncertainty.

[1] Warwick Manufacturing Group, University of Warwick, Coventry, UK. [2] Department of Computer Science, University of York, York, UK. [3] Intelligent and Autonomous Systems Group, Centrum Wiskunde & Informatica, Amsterdam, Netherlands. [4] Electrical Engineering Department, Eindhoven University of Technology, Eindhoven, Netherlands. [5] James Watt School of Engineering, University of Glasgow, Glasgow, UK. ✉email: xingyu.zhao@warwick.ac.uk; simos.gerasimou@york.ac.uk

Mobile robots are increasingly used to perform critical missions in extreme environments, which are inaccessible or hazardous to humans[1–4]. These missions range from the inspection and maintenance of offshore wind-turbine mooring chains and high-voltage cables to nuclear reactor repair and deep-space exploration[5,6].

Using robots for such missions poses major challenges[2,7]. First and foremost, the robots need to operate with high levels of autonomy, as in these harsh environments their interaction and communication with human operators is severely restricted. Additionally, they frequently need to make complex mission-critical decisions, with errors endangering not just the robot—itself an expensive asset, but also the important system or environment being inspected, repaired or explored. Last but not least, they need to cope with the considerable uncertainty associated with these missions, which often comprise one-off tasks or are carried out in settings not encountered before.

Addressing these major challenges is the focus of intense research worldwide. In the UK alone, a recent £44.5M research programme has tackled technical and certification challenges associated with the use of robotics and AI in the extreme environments encountered in offshore energy (https://orcahub.org), space exploration (https://www.fairspacehub.org), nuclear infrastructure (https://rainhub.org.uk), and management of nuclear waste (https://www.ncnr.org.uk). This research has initiated a step change in the assurance and certification of autonomous robots—not least through the emergence of new concepts such as dynamic assurance[8] and self-certification[9] for robotic systems.

Dynamic assurance requires a robot to respond to failures, environmental changes and other disruptions not only by reconfiguring accordingly[10], but also by producing new assurance evidence which guarantees that the reconfigured robot will continue to achieve its mission goals[8]. Self-certifying robots must continually verify their health and ability to complete missions in dynamic, risk-prone environments[9]. In line with the “defence in depth” safety engineering paradigm[11], this runtime verification has to be performed independently of the front-end planning and control engine of the robot.

Despite these advances, current dynamic assurance and self-certification methods rely on quantitative verification techniques (e.g., probabilistic[12,13] and statistical[14] model checking) that do not handle well the parametric uncertainty that autonomous robots encounter in extreme environments. Indeed, quantitative verification operates with stochastic models that demand single-point estimates of uncertain parameters such as task execution and failure rates. These estimates capture neither epistemic nor aleatory parametric uncertainty. As such, they are affected by arbitrary estimation errors which—because stochastic models are often nonlinear—can be amplified in the verification process[15], and may lead to invalid robot reconfiguration decisions, dynamic assurance and self-certification.

In this paper, we present a robust quantitative verification framework that employs Bayesian learning techniques to overcome this limitation. Our framework requires only partial and limited prior knowledge about the verified robotic system, and exploits its runtime observations (or lack thereof) to learn ranges of values for the system parameters. These parameter ranges are then used to compute the quantitative properties that underpin the robot’s decision making (e.g., probability of mission success, and expected energy usage) as intervals that—unique to our framework—capture the parametric uncertainty of the mission. Our framework is underpinned by probabilistic model checking, a technique that is broadly used to assess quantitative properties, e.g., reliability, performance and energy cost of systems exhibiting stochastic behaviour. Such systems include autonomous robots from numerous domains[16], e.g., mobile service robots[17], spacecraft[18], drones[19] and robotic swarms[20]. While we present a case study involving an autonomous underwater vehicle (AUV), the generalisability of our approach stems from the broad adoption of probabilistic model checking for the modelling and verification of this wide range of autonomous robots. As such, we anticipate that our results are applicable to autonomous agents across all these domains.

We start by introducing our robust verification framework, which comprises Bayesian techniques for learning the occurrence rates of both singular events (e.g., catastrophic failures and completion of one-off tasks) and events observed regularly during system operation. Next, we describe the use of the framework for an offshore wind turbine inspection and maintenance robotic mission. Finally, we discuss the framework in the context of related work, and we suggest directions for further research.

## Results

**Proposed framework**. We developed an end-to-end verification framework for the online computation of bounded intervals for continuous-time Markov chain (CMTC) properties that correspond to key dependability and performance properties of autonomous robots. The verification framework integrates interval CTMC model checking[21] with two new interval Bayesian inference techniques that we introduce in the Methods section. The former technique, Bayesian inference using partial priors (BIPP), computes estimate bounded intervals for the occurrence rates of singular events such as the successful completion of one-off robot tasks, or catastrophic failures. The latter technique, Bayesian inference using imprecise probability with sets of priors (IPSP), produces estimate bounded intervals for the occurrence rates of regular events encountered by an autonomous robot.

As shown in Fig. 1, the verification process underpinning the application of our framework involves devising a parametric CTMC model that captures the structural aspects of the system under verification through a SYSTEM MODELLER. This activity is typically performed once at design time (i.e., before the system is put into operation) by engineers with modelling expertise. By monitoring the system under verification after deployment, our framework enables observing both the occurrence of regular events and the lack of singular events during times when such events could have occurred (e.g., a catastrophic failure not happening when the system performs a dangerous operation). Our online BIPP ESTIMATOR and IPSP ESTIMATOR use these observations to calculate expected ranges for the rates of the monitored events, enabling a MODEL GENERATOR to continually synthesise up-to-date interval CTMCs that model the evolving behaviour of the system.

The interval CTMCs, which are synthesised from the parametric CTMC model, are then continually verified by the PRISM-PSY MODEL CHECKER[22], to compute value intervals for key system properties. As shown in Fig. 1 and illustrated in the next section, these properties range from dependability (e.g., safety, reliability and availability)[23] and performance (e.g., response time and throughput) properties to resource use and system utility. Finally, changes in the value ranges of these properties may prompt the dynamic reconfiguration of the system by a CONTROLLER module responsible for ensuring that the system requirements are satisfied at all times.

**Offshore infrastructure maintenance**. We demonstrate how our online robust verification and reconfiguration framework can support an AUV to execute a structural health inspection and cleaning mission of the substructure of an offshore wind farm. Similar scenarios for AUV use in remote, subsea environments have been described in other large-scale robotic demonstration

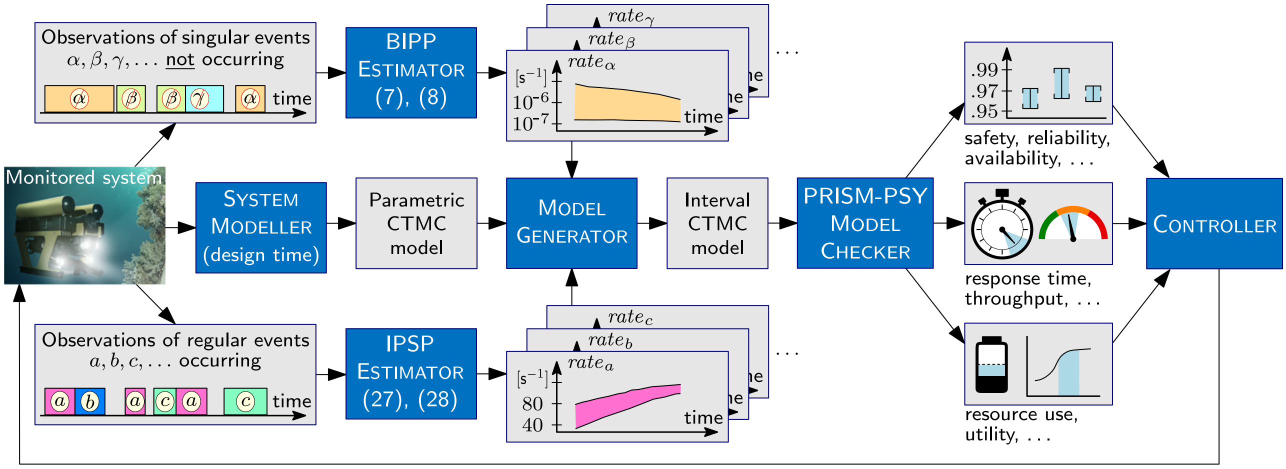

**Fig. 1 Robust Bayesian verification framework.** The integration of Bayesian inference using partial priors (BIPP) and Bayesian inference using imprecise probability with sets of priors (IPSP) with interval continuous-time Markov chain (CMTC) model checking supports the online robust quantitative verification and reconfiguration of autonomous systems under parametric uncertainty.

projects, such as the PANDORA EU FP7 project[24]. Compared to remotely operated vehicles that must be tethered with expensive oceanographic surface vessels run by specialised personnel, AUVs bring important advantages, including reduced environmental footprint (since no surface vessel consuming fuel is needed), reduced cognitive fatigue for the involved personnel, increased frequency of mission execution, and reduced operational and maintenance cost.

The offshore wind farm comprises multiple floating wind turbines, with each turbine being a buoyant foundation structure secured to the sea bed with floating chains tethered to anchors weighing several tons. Wind farms with floating wind turbines offer increased wind exploitation (since they can be installed in deeper waters where winds are stronger and more consistent), reduced installation costs (since there is no need to build solid foundations), and reduced impact on the visual and maritime life (since they are further from the shore)[25].

The AUV is deployed to collect data about the condition of $k \geq 1$ floating chains to enable the post-mission identification of problems that could affect the structural integrity of the asset (floating chain). When the visual inspection of a chain is hindered due to accumulated biofouling or marine growth, the AUV can use its on-board high-pressure water jet to clean the chain and continue with the inspection task[24].

The high degrees of *aleatoric uncertainty* in navigation and the perception of the marine environment entail that the AUV might fail to clean a chain. This uncertainty originates from the dynamic conditions of the underwater medium that includes unexpected water perturbations coupled with difficulties in scene understanding due to reduced visibility and the need to operate close to the floating chains. When this occurs, the AUV can retry the cleaning task or skip the chain and move to the next.

**Stochastic mission modelling**. Figure 2 shows the parametric CMTC model of the floating chain inspection and cleaning mission. The AUV inspects the *i*th chain with rate $r^{\text{inspect}}$ and consumes energy $e_{\text{ins}}$. The chain is clean with probability $p_c$ and the AUV travels to the next chain with rate $r^{\text{travel}}$ consuming energy $e_t$, or the chain needs cleaning with probability $1 - p_c$. When the AUV attempts the cleaning ($x_i = 1$), the task succeeds with chain-dependent rate $r_i^{\text{clean}}$, causes catastrophic damage to the floating chain or itself with rate $r^{\text{damage}}$ or fails with chain-dependent rate $r_i^{\text{fail}}$. If the cleaning fails, the AUV prepares to retry with known and fixed rate $r^{\text{prepare}}$ requiring energy $e_p$, and it either retries cleaning ($x_i = 1$) or skips the current chain and moves to chain $i + 1$ ($x_i = 0$). After executing the tasks on the *k*th chain, the AUV returns to its base and completes the mission.

Since the AUV can fail to clean the *i*-th chain with non-negligible probability and multiple times, this is a regular event whose transition rate $r_i^{\text{fail}}$ is modelled using the IPSP estimator from (7) and (8). In contrast, the AUV is expected to not cause catastrophic damage but, with extremely low probability, may do so only once (after which the AUV and/or its mission are likely to be revised); thus, the corresponding transition rates $r_i^{\text{clean}}$ and $r^{\text{damage}}$ are modelled using the BIPP estimator from (14) and (15). The other transition rates, i.e., those for inspection ($r^{\text{inspect}}$), travelling ($r^{\text{travel}}$) and preparation ($r^{\text{prepare}}$), are less influenced by the chain conditions and therefore assumed to be known, e.g., from previous trials and missions; hence, we fixed these transition rates.

When cleaning is needed for the *i*th chain, the AUV controller synthesises a plan by determining the control parameter $x_i \in \{0, 1\}$ for all remaining chains $i, i+1, \ldots k$ so that the system requirements in Table 1 are satisfied.

**Robust verification results**. We demonstrate our solution for robust verification and adaptation using a mission in which the AUV was deployed to inspect and, if needed, clean six chains placed in a hexagonal arrangement (Fig. 3). We used $m = 3$ and the BIPP estimator (7) and (8) for the transition rates $r_i^{\text{clean}}$ and $r^{\text{damage}}$, which correspond to singular events. For $r_i^{\text{clean}}$, we used $\epsilon_1 = 0.12 + \mathcal{U}(0, 0.12)$, $\theta_1 = 0.10 + \mathcal{U}(0, 0.001)$, $\epsilon_2 = 0.90 + \mathcal{U}(0, 0.90)$, $\theta_2 = 0.85 + \mathcal{U}(0, 0.0085)$, where $\mathcal{U}(x, y)$ denotes a continuous uniform distribution with $x$ and $y$ being its minimum and maximum values, respectively. For $r^{\text{damage}}$, we used $\epsilon_1 = 10^{-8} + \mathcal{U}(0, 10^{-8})$, $\theta_1 = 0.88 + \mathcal{U}(0, 0.0088)$, $\epsilon_2 = 10^{-7} + \mathcal{U}(0, 10^{-7})$, $\theta_2 = 0.10 + \mathcal{U}(0, 0.001)$. For $r_i^{\text{fail}}$, we used $t^{(0)} = [10 + \mathcal{U}(0, 10)]$ and $\lambda^{(0)} = [0.0163 + \mathcal{U}(0, 0.00163)]$. During the mission execution, the AUV performs the model checking at every cleaning attempt so that runtime observations are incorporated into the decision making process entailing also that the currently synthesised plan is not necessarily used at subsequent chains. Hence, the AUV only needs to check system configurations where at least the current chain is to be cleaned, thus halving the number of configurations to be checked (since configurations with $x_i = 0$ need not be checked). If all of these

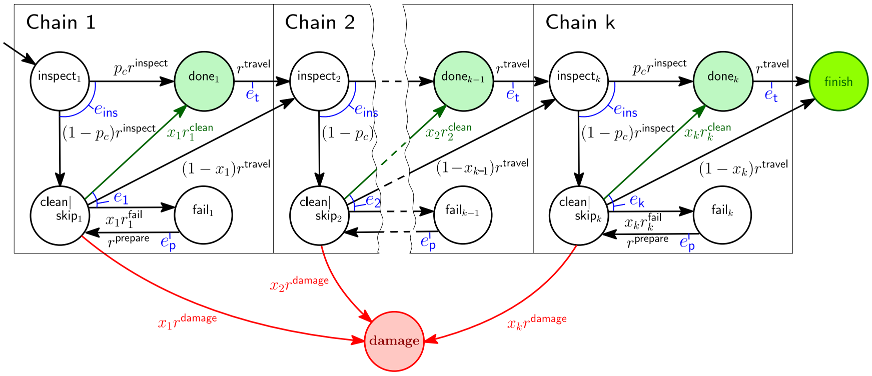

**Fig. 2 Floating chain continuous-time Markov chain (CMTC) model of an autonomous underwater vehicle (AUV).** CTMC of the floating chain cleaning and inspection mission, where $e_1$, $e_2$, …, $e_k$ represent the mean energy required to clean chains 1, 2, …, $k$, respectively. The AUV inspects a chain with rate $r^{\text{inspect}}$, consuming energy $e_{\text{ins}}$, prepares to retry the chain cleaning task with rate $r^{\text{prepare}}$, consuming energy $e_p$, and travels to the next chain with rate $r^{\text{travel}}$, consuming energy $e_t$. During an inspection, the chain is clean with probability $p_c$, and $x_1, x_2, …, x_k \in \{0, 1\}$ denote the control parameters used by the AUV controller to synthesise a plan. The rate $r_i^{\text{fail}}$ corresponds to a regular event and is therefore modelled using Bayesian inference using imprecise probability with sets of priors (IPSP) from (27) and (28). The rates $r_i^{\text{clean}}$ and $r^{\text{damage}}$ correspond to singular events and are thus modelled using Bayesian inference using partial priors (BIPP) from (7) and (8).

**Table 1 System requirements for the AUV floating chain inspection and cleaning mission**

| ID | Informal description | Formal specification[a] |
|---|---|---|
| R1 | The probability of mission failure must not exceed 5% | $P_{\leq 0.05}$ [$F$ damage] |
| R2 | The expected energy consumption must not exceed the remaining energy $E_{\text{left}}$ | $R^{\text{energy}}_{\leq E_{\text{left}}}$[$F$ finish] |
| R3 | Subject to R1 and R2 being met, maximise the number of cleaned chains | Find argmax $\sum_{i=1}^{k} x_i$ such that $R1 \wedge R2$ |

[a] Expressed in rewards-extended continuous stochastic logic (see Methods section).

checks that consider $x_i = 1$ fail to satisfy the requirements from Table 1, then the AUV decides to skip the current chain and proceed to inspect and clean the next chain.

If a cleaning attempt at chain $i$ failed, the AUV integrates this observation in (27) and (28), and performs model checking to determine whether to retry the cleaning or skip the chain. Since the AUV has consumed energy for the failed cleaning attempt, the energy available is reduced accordingly, which in turn can reduce the number of possible system configurations that can be employed and need checking. The observation of a failed attempt reduces the lower bound for the reliability of cleaning $x_i$, and may result in a violation of the reliability requirement R1 (Table 1), which may further reduce the number of feasible configurations. If the AUV fails to clean chain $i$ repeatedly, this lower bound will continue to decrease, potentially resulting in the AUV having no feasible configuration, and having to skip the current chain. Although skipping a chain overall decreases the risk of a catastrophic failure (as the number of cleaning attempts is reduced), leaving uncleaned chains will incur additional cost as a new inspection mission will need to be launched, e.g., using another AUV or human personnel.

Figure 3 shows a simulated run of the AUV performing an inspection and cleaning mission (Fig. 3a). At each chain that requires cleaning, the AUV decides whether to attempt to clean or skip the current chain. Figure 3b provides details of the probabilistic model checking carried out during the inspection and cleaning of chain 3 (Fig. 3a, ii). Overall, the AUV performed multiple attempts to clean chain 3, succeeding on the third attempt.

The results of the model checking analyses for these attempts are shown in successive columns in Fig. 3b, while each row depicts the analysis of one of the requirements from Table 1. A system configuration is feasible if it satisfies requirements R1—the AUV will not encounter a catastrophic failure with a probability of at least 0.95, and R2—the expected energy consumption does not exceed the remaining AUV energy. Lastly, if multiple configurations satisfy requirements R1 and R2, then the winner is the configuration that maximises the number of chains cleaned. If there is still a tie, the configuration is chosen randomly from those that clean the most chains.

In the AUV's first attempt at chain 3 (Fig. 3b (i–iii)), all the configurations are feasible, so configuration 1 (highlighted, and corresponding to the highest number of chains cleaned) is selected. This attempt fails, and a second assessment is made (Fig. 3b (iv–vi)). This time, only system configurations 2–8 are feasible, and as configurations 2, 3, and 5 maximise R3, a configuration is chosen randomly from this subset (in this case, configuration 3). This attempt also fails, and on the third attempt (Fig. 3b (vii–ix)), only configurations 4–8 are feasible, with 5 maximising R3, and the AUV adopts this configuration and succeeds in cleaning the chain.

In this AUV mission instance, the AUV controller is concerned with cleaning the maximum number of chains and ensuring the AUV returns safely. In other variants of our AUV mission, the system properties from requirements R1 and R2 could also be used to determine a winning configuration in the event of a tie between multiple feasible configurations. For example, it might be optimal for the AUV to consume minimal energy in this scenario. Thus, the energy consumption from requirement R2 can be used as a metric to choose a configuration as a tie-breaker.

We also measured the overheads associated with executing the online verification process. Figure 4 shows the computation overheads incurred by the RBV framework for executing the AUV-based mission. The values comprising each boxplot have been collected over 10 independent runs. Each value denotes the time consumed for a single online robust quantitative

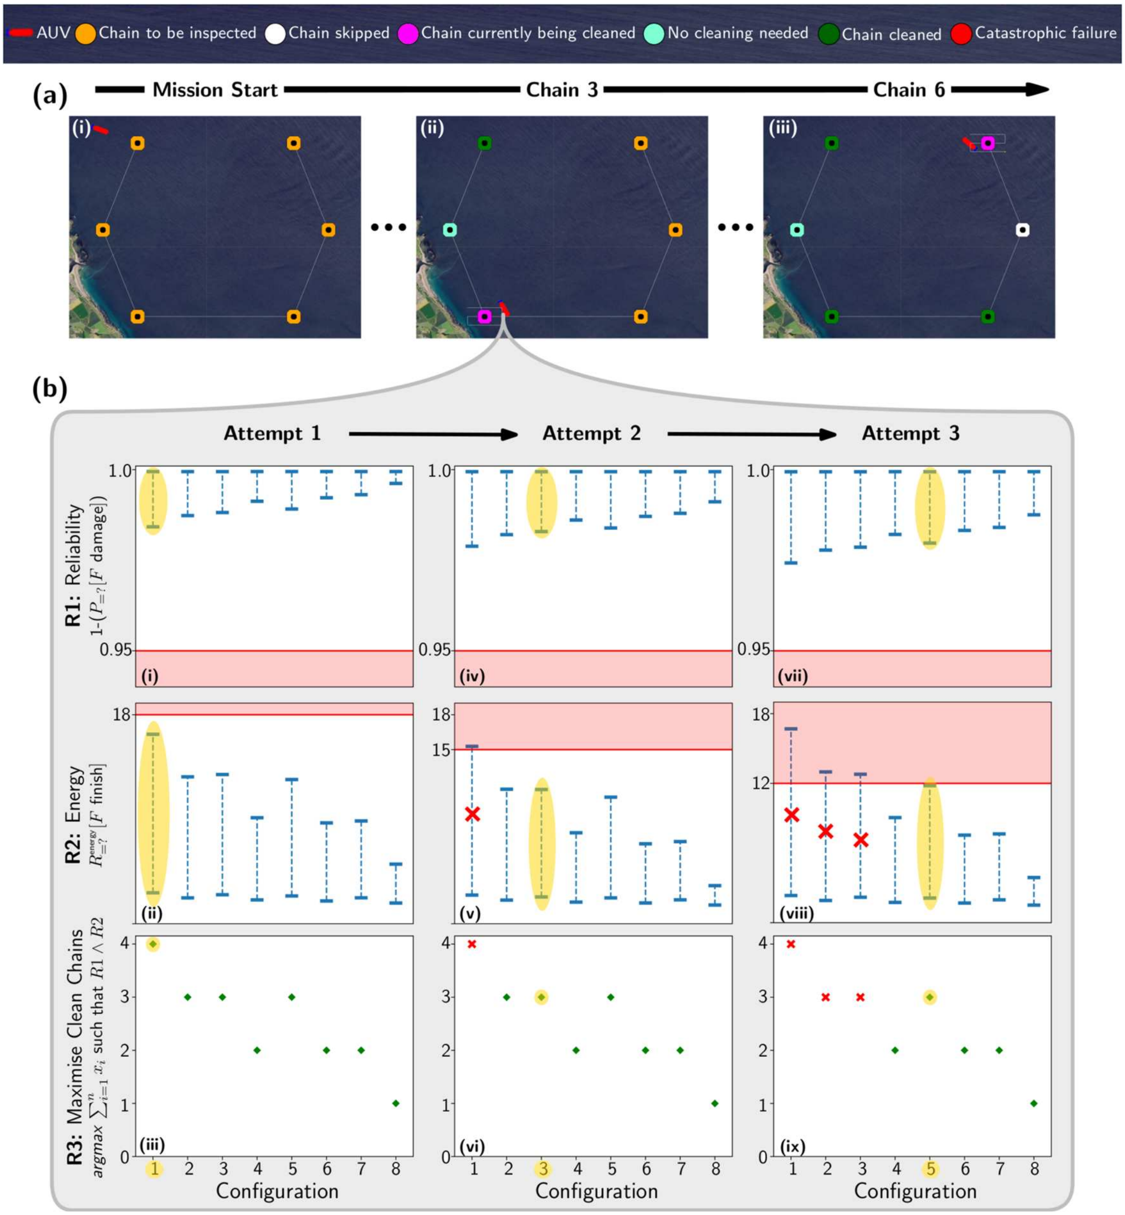

**Fig. 3 Demonstration of autonomous underwater vehicle (AUV) inspection and cleaning mission. a** Simulated AUV mission involving the inspection of six wind farm chains and, if required, their cleaning. **i** Start of mission; **ii** cleaning chain 3; **iii** cleaning final chain. At this point, the AUV cleaned three chains, skipped one, and one chain did not require cleaning. **b** Plots of the outcome of the model checking carried out by the AUV at chain 3. Each row shows the configurations against the requirements. **i**–**iii** Results during the first attempt at cleaning chain 3. **iv**–**vi** Results during the second attempt at cleaning. **vii**–**ix** Results at the third and successful attempt at cleaning the chain. The configurations decorated with the red cross signify configurations violating the energy requirement R2 while configurations highlighted in yellow denote the chosen configuration for the corresponding attempt.

verification and reconfiguration step when the AUV attempts to clean the indicated chain. For instance, the boxplot associated with the 'Chain 1' ('Chain 2') label on the x-axis signifies that the AUV attempts to clean chain 1 (chain 2) and corresponds to the time consumed by the RBV framework to analyse 64 (32) configurations. Overall, the time overheads are reasonable for the purpose of this mission. Since the AUV has more configurations to analyse at the earlier stages of the mission (e.g., when inspecting chain 1), the results follow the anticipated exponential pattern. The number of configurations decreases by half each time the AUV progresses further into the mission and moves to the next chain. Another interesting observation is that the length of each boxplot is small, i.e., the lower and upper quartiles are very close, indicating that the RBV framework showcases a consistent behaviour in the time taken for its execution.

The consumed time comprises (1) the time required to compute the posterior estimate bounds of the modelled transition rates, $r_i^{\text{clean}}$, $r^{\text{fail}}$, $1 \le i \le k$, and $r^{\text{damage}}$, using the BIPP and IPSP estimators; (2) the time required to compute the value intervals for requirements R1 and R2 using the probabilistic model checker PRISM-PSY[22]; and (3) the time needed to find the best configuration satisfying requirements R1 and R2, and maximising requirement R3. Our empirical analysis provided evidence that the execution of the BIPP and IPSP estimators and the selection of the best configuration have negligible overheads with almost all time incurred by PRISM-PSY. This outcome is not surprising and is

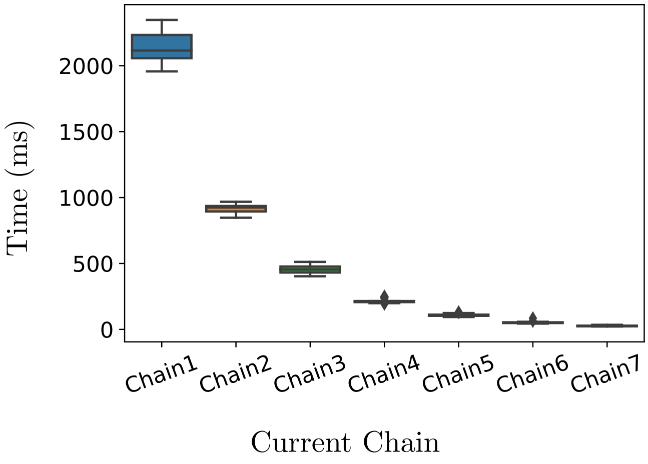

**Fig. 4 Verification time overheads.** Time taken by our robust Bayesian verification framework to execute the online quantitative verification and reconfiguration step over 10 independent runs when the robot attempts to clean the indicated chain.

aligned with the results reported in[22] concerning the execution overheads of the model checker.

Additional information about the offshore infrastructure maintenance experiments, including details about the experimental methodology, is provided in Supplementary Methods 2 of the Supplementary material. The simulator used for the AUV mission, developed on top of the open-source MOOS-IvP middleware[26], and a video showing the execution of this AUV mission instance are available at http://github.com/gerasimou/RBV.

## Discussion

Unlike single-point estimators of Markov model parameters[27–30], our Bayesian framework provides interval estimates that capture the inherent uncertainty of these parameters, enabling the robust quantitative verification of systems such as autonomous robots. Through its ability to exploit prior knowledge, the framework differs fundamentally from, and is superior to, a recently introduced approach to synthesising intervals for unknown transition parameters based on the frequentist theory of simultaneous confidence intervals[15,31,32]. Furthermore, instead of applying the same estimator to all Markov model transition parameters like existing approaches, our framework is the first to handle parameters corresponding to singular and regular events differently. This is an essential distinction, especially for the former type of parameter, for which the absence of observations violates a key premise of existing estimators. Our BIPP estimator avoids this invalid premise, and computes two-sided bounded estimates for singular CTMC transition rates—a considerable extension of our preliminary work to devise one-sided bounded estimates for the singular transition probabilities of discrete-time Markov chains[33].

The proposed Bayesian framework is underpinned by the theoretical foundations of imprecise probabilities[34,35] and Conservative Bayesian Inference (CBI)[36–38] integrated with recent advances in the verification of interval CTMCs[22]. In particular, our BIPP theorems for singular events extend CBI significantly in several ways. First, BIPP operates in the continuous domain for a Poisson process, while previous CBI theorems are applicable to Bernoulli processes in the discrete domain. As such, BIPP enables the runtime quantitative verification of interval CTMCs, and thus the analysis of important properties that are not captured by discrete-time Markov models. Second, CBI is one-side (upper) bounded, and therefore only supports the analysis of undesirable singular events (e.g., catastrophic failures). In contrast, BIPP provides two-sided bounded estimates, therefore also enabling the analysis of "positive" singular events (e.g., the completion of difficult one-off tasks). Finally, BIPP can operate with any *arbitrary* number of confidence bounds as priors, which greatly increases the flexibility of exploiting different types of prior knowledge.

As illustrated by its application to an AUV infrastructure maintenance mission, our robust quantitative verification framework removes the need for precise prior beliefs, which are typically unavailable in many real-world verification tasks that require Bayesian inference. Instead, the framework enables the exploitation of Bayesian combinations of partial or imperfect prior knowledge, which it uses to derive informed estimation errors (i.e., intervals) for the predicted model parameters. Combined with existing techniques for obtaining this prior knowledge, e.g., the Delphi method and its variants[39] or reference class forecasting[40], the framework increases the trustworthiness of Bayesian inference in highly uncertain scenarios such as those encountered in the verification of autonomous robots.

Based on recent survey papers[41–43] that provide in-depth discussions on the challenges and opportunities in the field of autonomous robot verification, it has become evident that a common taxonomy emerges, primarily revolving around two key dimensions. The first dimension centres on the specification of properties under verification, which includes various types of temporal logic languages[41]. The second dimension pertains to how system behaviours are modeled/structured. In this regard, formal models such as Belief Desire Intention, Petri Nets, and finite state machines, along with their diverse extensions, have emerged as popular approaches to capturing the intricate dynamics of autonomous systems. Our approach falls within the category of methods utilising continuous stochastic logic (CSL) and CTMCs for the verification of robots. However, unlike the existing methods from this category[44,45], we introduced treatments of the model parameters uncertainty via robust Bayesian learning methods, and integrated them with recent research on interval CMTC model checking.

Another important approach for verifying the behaviour of autonomous agents under uncertainty uses hidden Markov models (HMMs)[46–48]. HMM-based verification supports the analysis of stochastic systems whose true state is not observable, and can only be estimated (with aleatoric uncertainty given by a predefined probability distribution) through monitoring a separate process whose observable state depends on the unknown state of the system. In contrast, our verification framework supports the analysis of autonomous agents whose true state is observable but for which the rates of transition between these known states are affected by epistemic uncertainty and need to be learnt from system observations (as shown in Fig. 1). As such, HMM-based verification and our robust verification framework differ substantially by tackling different types of autonomous agent uncertainty. Because autonomous agents may be affected by both types of uncertainty, the complementarity of the two verification approaches can actually be leveraged by using our BIPP and IPSP Bayesian estimators in conjunction with HMM-based verification, i.e., to learn the transition rates associated with continuous-time HMMs that model the behaviour of an autonomous agent. Nevertheless, achieving this integration will first require the development of generic continuous-time HMM verification techniques since, to the best of our knowledge, only verification techniques and tools for the verification of discrete-time HMMs are currently available.

Although our method demonstrates promising potential, it is not without limitations. One limitation is scalability—as the complexity of the robot's behaviour and the environment grow, the number of unknown parameters to be estimated at runtime may increase, leading to increased computational overheads for our Bayesian estimators. Additionally, the method requires a certain level of expertise to construct the underlying CTMC model structure. This demands understanding both the robot's dynamics and the environment in order to model them as a CTMC, making the approach less accessible to those without specialised knowledge. Last but not least, a challenge inherent to all Bayesian methods involves the acquisition of appropriate priors. While our robust Bayesian estimators mitigate this issue by eliminating the need for complete and precise prior knowledge, establishing the required partial and vague priors can still pose challenges. These limitations suggest important areas for future work.

## Methods

**Quantitative verification**. Quantitative verification is a mathematically based technique for analysing the correctness, reliability, performance and other key properties of systems with stochastic behaviour[49,50]. The technique captures this behaviour into Markov models, formalises the properties of interest as probabilistic temporal logic formulae over these models, and employs efficient algorithms for their analysis. Examples of such properties include the probability of mission failure for an autonomous robot, and the expected battery energy required to complete a robotic mission.

In this paper, we focus on the quantitative verification of continuous-time Markov chains (CMTCs). CTMCs are Markov models for continuous-time stochastic processes over countable state spaces comprising (i) a finite set of states corresponding to real-world states of the system that are relevant for the analysed properties; and (ii) the rates of transition between these states. We use the following definition adapted from the probabilistic model checking literature[49,50].

**Definition 1**. A continuous-time Markov chain is a tuple

$$\mathcal{M} = (S, s_0, \mathbf{R}), \tag{1}$$

where $S$ is a finite set of states, $s_0 \in S$ is the initial state, and $\mathbf{R} : S \times S \rightarrow \mathbb{R}$ is a transition rate matrix such that the probability that the CTMC will leave state $s_i \in S$ within $t > 0$ time units is $1 - e^{-t \cdot \sum_{s_k \in S \setminus \{s_i\}} \mathbf{R}(s_i, s_k)}$ and the probability that the new state is $s_j \in S \setminus \{s_i\}$ is $p_{ij} = \mathbf{R}(s_i, s_j) \,/\, \sum_{s_k \in S \setminus \{s_i\}} \mathbf{R}(s_i, s_k)$.

The range of properties that can be verified using CTMCs can be extended by annotating the states and transitions with non-negative quantities called rewards.

**Definition 2**. A reward structure over a CTMC $\mathcal{M} = (S, s_0, \mathbf{R})$ is a pair of functions $(\underline{\rho}, \boldsymbol{\iota})$ such that $\underline{\rho} : S \rightarrow \mathbb{R}_{\geq 0}$ is a state reward function (a vector), and $\boldsymbol{\iota} : S \times S \rightarrow \mathbb{R}_{\geq 0}$ is a transition reward function (a matrix).

CTMCs support the verification of quantitative properties expressed in CSL[51] extended with rewards[50].

**Definition 3**. Given a set of atomic propositions $AP$, $a \in AP$, $p \in [0, 1]$, $I \subseteq \mathbb{R}_{\geq 0}$, $r, t \in \mathbb{R}_{\geq 0}$ and $\bowtie \;\in \{\geq, >, <, \leq\}$, a CSL formula $\Phi$ is defined by the grammar:

$$\Phi ::= true \mid a \mid \Phi \wedge \Phi \mid \neg\Phi \mid P_{\bowtie p}[X\,\Phi] \mid P_{\bowtie p}[\Phi\ U^I\ \Phi] \mid \mathcal{S}_{\bowtie p}[\Phi] \mid$$

$$R_{\bowtie r}[I^{=t}] \mid R_{\bowtie r}[C^{\leq t}] \mid R_{\bowtie r}[F\,\Phi] \mid R_{\bowtie r}[S].$$

Given a CTMC $\mathcal{M} = (S, s_0, \mathbf{R})$ with states labelled with atomic propositions from $AP$ by a function $L : S \rightarrow 2^{AP}$, and a reward structure $(\underline{\rho}, \boldsymbol{\iota})$ over $\mathcal{M}$, the CSL semantics is defined with a satisfaction relation $\models$ over the states and paths (i.e., feasible sequences of successive states) of $\mathcal{M}$[49]. The notation $s \models \Phi$ means "$\Phi$ is satisfied in state $s$". For any state $s \in S$, we have:

- $s \models true$, $s \models a$ iff $a \in L(s)$, $s \models \neg\Phi$ iff $\neg(s \models \Phi)$, and $s \models \Phi_1 \wedge \Phi_2$ iff $s \models \Phi_1$ and $s \models \Phi_2$;
- $s \models \mathcal{P}_{\bowtie p}[X\,\Phi]$ iff the probability $x$ that $\Phi$ holds in the state following $s$ satisfies $x \bowtie p$ (probabilistic next formula);
- $s \models \mathcal{P}_{\bowtie p}[\Phi_1\ U^I\ \Phi_2]$ iff, across all paths starting at $s$, the probability $x$ of going through only states where $\Phi_1$ holds until reaching a state where $\Phi_2$ holds at a time $t \in I$ satisfies $x \bowtie p$ (probabilistic until formula);
- $s \models \mathcal{S}_{\bowtie p}[\Phi]$ iff, having started in state $s$, the probability $x$ of $\mathcal{M}$ reaching a state where $\Phi$ holds in the long run satisfies $x \bowtie p$ (probabilistic steady-state formula);
- the instantaneous $(R_{\bowtie r}[I^{=t}])$, cumulative $(R_{\bowtie r}[C^{\leq t}])$, future-state $(R_{\bowtie r}[F\,\Phi])$ and steady-state $(R_{\bowtie r}[S])$ reward formulae hold iff, having started in state $s$, the expected reward $x$ at time instant $t$, cumulated up to time $t$, cumulated until reaching a state where $\Phi$ holds, and achieved at steady state, respectively, satisfies $x \bowtie r$.

Probabilistic model checkers such as PRISM[52] and Storm[53] use efficient analysis techniques to compute the actual probabilities and expected rewards associated with probabilistic and reward formulae, respectively. The formulae are then verified by comparing the computed values to the bounds $p$ and $r$. Furthermore, the extended CSL syntax $P_{=?}[X\,\Phi]$, $P_{=?}[\Phi_1\ U^I\ \Phi_2]$, $R_{=?}[I^{=t}]$, etc. can be used to obtain these values from the model checkers.

While the transition rates of the CTMCs verified in this way must be known and constant, advanced quantitative verification techniques[21] support the analysis of CTMCs whose transition rates are specified as intervals. The technique has been used to synthesise CTMCs corresponding to process configurations and system designs that satisfy quantitative constraints and optimisation criteria[22,32,54], under the assumption that these bounded intervals are available. Here we introduce a Bayesian framework for computing these intervals in ways that reflect the parametric uncertainty of real-world systems such as autonomous robots.

**Bayesian learning of CTMC transition rates**. Given two states $s_i$ and $s_j$ of a CTMC such that transitions from $s_i$ to $s_j$ are possible and occur with rate $\lambda$, each transition from $s_i$ to $s_j$ is independent of how state $s_i$ was reached (the Markov property). Furthermore, the time spent in state $s_i$ before a transition to $s_j$ is modelled by a homogeneous Poisson process of rate $\lambda$. Accordingly, the likelihood that 'data' collected by observing the CTMC shows $n$ such transitions occurring within a combined time $t$ spent in state $s_i$ is given by the conditional probability:

$$l(\lambda) = Pr(\text{data}|\lambda) = \frac{(\lambda t)^n}{n!} e^{-\lambda t} \tag{2}$$

In practice, the rate $\lambda$ is typically unknown, but prior beliefs about its value are available (e.g., from domain experts or from past missions performed by the system modelled by the CTMC) in the form of a probability (density or mass) function $f(\lambda)$. In this common scenario, the Bayes Theorem can be used to derive a posterior probability function that combines the likelihood $l(\lambda)$ and the prior $f(\lambda)$ into a better estimate for $\lambda$ at time $t$:

$$f(\lambda|\text{data}) = \frac{l(\lambda) f(\lambda)}{\int_0^\infty l(\lambda) f(\lambda) \mathrm{d}\lambda} \tag{3}$$

where the Lebesgue-Stieltjes integral from the denominator is introduced to ensure that $f(\lambda|\text{data})$ is a probability function. We note, we use Lebesgue-Stieltjes integration to cover in a compact way both continuous and discrete prior distributions $f(\lambda)$, as these integrals naturally reduce to sums for discrete distributions. We calculate the posterior estimate for the rate $\lambda$ at time $t$ as the expectation of (3):

$$\lambda^{(t)} = \mathbb{E}[\Lambda|\text{data}] = \frac{\int_0^\infty \lambda l(\lambda) f(\lambda) \mathrm{d}\lambda}{\int_0^\infty l(\lambda) f(\lambda) \mathrm{d}\lambda}. \tag{4}$$

where we use capital letters for random variables and lower case for their realisations.

**Interval Bayesian inference for singular events**. In the autonomous-robot missions considered in our paper, certain events are extremely rare, and treated as unique from a modelling viewpoint. These events include major failures (after each of which the system is modified to remove or mitigate the cause of the failure), and the successful completion of difficult one-off tasks. Using Bayesian inference to estimate the CTMC transition rates associated with such events is challenging because, with no observations of these events, the posterior estimate is highly sensitive to the choice of a suitable prior distribution. Furthermore, only limited domain knowledge is often available to select and justify a prior distribution for these singular events.

To address this challenge, we develop a Bayesian inference using partial priors (BIPP) estimator that requires only *limited, partial prior knowledge* instead of the complete prior distribution typically needed for Bayesian inference. For one-off events, such knowledge is both more likely to be available and easier to justify. BIPP provides bounded posterior estimates that are robust in the sense that the ground truth rate values are within the estimated intervals.

To derive the BIPP estimator, we note that for one-off events the likelihood (2) becomes

$$l(\lambda) = Pr(\text{data}|\lambda) = e^{-\lambda \cdot t} \tag{5}$$

because $n = 0$. Instead of a prior distribution $f(\lambda)$ (required to compute the posterior expectation (4)), we assume that we only have limited partial knowledge consisting of $m \geq 2$ confidence bounds on $f(\lambda)$:

$$Pr(\epsilon_{i-1} < \lambda \leq \epsilon_i) = \theta_i \tag{6}$$

where $1 \leq i \leq m$, $\theta_i > 0$, and $\sum_{i=1}^m \theta_i = 1$. The use of such bounds is a common practice for safety-critical systems. As an example, the IEC61508 safety standard [55] defines safety integrity levels (SILs) for the critical functions of a system based on the bounds for their probability of failure on demand (*pfd*): *pfd* between $10^{-2}$ and $10^{-1}$ corresponds to SIL 1, *pfd* between $10^{-3}$ and $10^{-2}$ corresponds to SIL 2, etc.; and testing can be used to estimate the probabilities that a critical function has different SILs. We note that $Pr(\lambda \geq \epsilon_0) = Pr(\lambda \leq \epsilon_m) = 1$ and that, when no specific information is available, we can use $\epsilon_0 = 0$ and $\epsilon_m = +\infty$.

The partial knowledge encoded by the constraints (6) is far from a complete prior distribution: an infinite number of distributions $f(\lambda)$ satisfy these constraints, and the result below provides bounds for the estimate rate (4) across these distributions.

**Theorem 1**. The set $S_\lambda$ of posterior estimate rates (4) computed for all prior distributions $f(\lambda)$ that satisfy (6) has an infinum $\lambda_l$ and a supremum $\lambda_u$ given by:

$$\lambda_l = \min\left\{ \frac{\sum_{i=1}^m \left[\epsilon_i l(\epsilon_i)(1-x_i)\theta_i + \epsilon_{i-1} l(\epsilon_{i-1}) x_i \theta_i\right]}{\sum_{i=1}^m [l(\epsilon_i)(1-x_i)\theta_i + l(\epsilon_{i-1}) x_i \theta_i]} \middle| \forall 1 \leq i \leq m . x_i \in [0,1] \right\}, \tag{7}$$

$$\lambda_u = \max\left\{ \frac{\sum_{i=1}^m \lambda_i l(\lambda_i)\theta_i}{\sum_{i=1}^m l(\lambda_i)\theta_i} \middle| \forall 1 \leq i \leq m . \lambda_i \in (\epsilon_{i-1}, \epsilon_i] \right\}. \tag{8}$$

Before providing a proof for Theorem 1, we note that the values $\lambda_l$ and $\lambda_u$ can be computed using numerical optimisation software packages available, for instance, within widely used mathematical computing tools like MATLAB and Maple. For applications where computational resources are limited or the BIPP estimator is used online with tight deadlines, the following corollaries (whose proofs are provided in our supplementary material) give closed-form estimator bounds for $m = 3$ (with $m = 2$ as a subcase).

**Corollary 1**. When $m = 3$, the bounds (7) and (8) satisfy:

$$\lambda_l \geq \begin{cases} \frac{\epsilon_1 l(\epsilon_1)\theta_2}{\theta_1 + l(\epsilon_1)\theta_2}, & \text{if } \frac{\theta_2(\epsilon_1 - \epsilon_2)}{\theta_1} > \frac{\epsilon_2 l(\epsilon_2) - \epsilon_1 l(\epsilon_1)}{l(\epsilon_1) l(\epsilon_2)} \\ \frac{\epsilon_2 l(\epsilon_2)\theta_2}{\theta_1 + l(\epsilon_2)\theta_2}, & \text{otherwise} \end{cases} \tag{9}$$

and

$$\lambda_u < \begin{cases} \frac{\epsilon_1 l(\epsilon_1)\theta_1 + \epsilon_2 l(\epsilon_2)\theta_2 + \frac{1}{t} l(\frac{1}{t})(1-\theta_1-\theta_2)}{l(\epsilon_1)\theta_1}, & \text{if } t < \frac{1}{\epsilon_2} \\ \frac{\epsilon_1 l(\epsilon_1)\theta_1 + \frac{1}{t} l(\frac{1}{t})\theta_2 + \epsilon_2 l(\epsilon_2)(1-\theta_1-\theta_2)}{l(\epsilon_1)\theta_1}, & \text{if } \frac{1}{\epsilon_2} \leq t \leq \frac{1}{\epsilon_1} \\ \frac{\epsilon_1 l(\epsilon_1)(\theta_1+\theta_2) + \epsilon_2 l(\epsilon_2)(1-\theta_1-\theta_2)}{l(\epsilon_1)\theta_1}, & \text{otherwise} \end{cases} \tag{10}$$

**Corollary 2**. Closed-form BIPP bounds for $m = 2$ can be obtained by setting $\epsilon_2 = \epsilon_1$ and $\theta_2 = 0$ in (9) and (10).

To prove Theorem 1, we require the following Lemma and Propositions.

**Lemma 1**. If $l(\cdot)$ is the likelihood function defined in (5), then $g : (0, \infty) \to \mathbb{R}$, $g(w) = w \cdot l^{-1}(w)$ is a concave function.

**Proof**. Since $g(w) = w \cdot \left(-\frac{\ln w}{t}\right)$ and $t > 0$, the second derivative of $g$ satisfies

$$\frac{d^2 g}{dw^2} = \frac{d}{dw}\left[-\frac{\ln w}{t} - \frac{1}{t}\right] = -\frac{1}{wt} < 0. \tag{11}$$

Thus, $g(w)$ is concave. □

**Proposition 1**. With the notation from Theorem 1, there exist $m$ values $\lambda_1 \in (\epsilon_0, \epsilon_1]$, $\lambda_2 \in (\epsilon_1, \epsilon_2]$, …, $\lambda_m \in (\epsilon_{m-1}, \epsilon_m]$ such that $\sup \mathcal{S}_\lambda$ is the posterior estimate (4) obtained by using as prior the $m$-point discrete distribution with probability mass $f(\lambda_i) = Pr(\lambda = \lambda_i) = \theta_i$ for $i = 1, 2, \ldots, m$.

**Proof**. Since $f(\lambda) = 0$ for $\lambda \notin [\epsilon_0, \epsilon_m]$, the Lebesgue-Stieltjes integration from the objective function (4) can be rewritten as:

$$\mathbb{E}(\Lambda|\text{data}) = \frac{\sum_{i=1}^m \int_{\epsilon_{i-1}}^{\epsilon_i} \lambda l(\lambda) f(\lambda) \mathrm{d}\lambda}{\sum_{i=1}^m \int_{\epsilon_{i-1}}^{\epsilon_i} l(\lambda) f(\lambda) \mathrm{d}\lambda} \tag{12}$$

The first mean value theorem for integrals (e.g.,[56] p. 249]) ensures that, for every $i = 1, 2, \ldots, m$, there are points $\lambda_i, \lambda'_i \in [\epsilon_{i-1}, \epsilon_i]$ such that:

$$\int_{\epsilon_{i-1}}^{\epsilon_i} l(\lambda) f(\lambda) \mathrm{d}\lambda = l(\lambda_i) \int_{\epsilon_{i-1}}^{\epsilon_i} f(\lambda) \mathrm{d}\lambda = l(\lambda_i)\theta_i \tag{13}$$

$$\int_{\epsilon_{i-1}}^{\epsilon_i} \lambda l(\lambda) f(\lambda) \mathrm{d}\lambda = \lambda'_i l(\lambda'_i) \int_{\epsilon_{i-1}}^{\epsilon_i} f(\lambda) \mathrm{d}\lambda = \lambda'_i l(\lambda'_i)\theta_i \tag{14}$$

or, after simple algebraic manipulations of the previous results,

$$l(\lambda_i) = \mathbb{E}[l(\Lambda)|\epsilon_{i-1} \leq \Lambda \leq \epsilon_i] \tag{15}$$

$$\lambda'_i l(\lambda'_i) = \mathbb{E}[\Lambda \cdot l(\Lambda)|\epsilon_{i-1} \leq \Lambda \leq \epsilon_i] \tag{16}$$

Using the shorthand notation $w = l(\lambda)$ for the likelihood function (5) (hence $w > 0$), we define $g : (0, \infty) \to \mathbb{R}$, $g(w) = w \cdot l^{-1}(w)$. According to Lemma 1, $g(\cdot)$ is a concave function, and thus we have:

$$
\begin{aligned}
\lambda_i' l(\lambda_i') &= \mathbb{E}[\Lambda \cdot l(\Lambda) | \epsilon_{i-1} \le \Lambda \le \epsilon_i] \\
&= \mathbb{E}[W \cdot l^{-1}(W) | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i] \\
&= \mathbb{E}[g(W) | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i] \\
&\le g\left(\mathbb{E}[W | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i]\right) \qquad (17)
\end{aligned}
$$

$$
\begin{aligned}
&= \mathbb{E}[W | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i] \cdot \\
&\quad l^{-1}\left(\mathbb{E}[W | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i]\right) \\
&= \mathbb{E}[l(\Lambda) | \epsilon_{i-1} \le \Lambda \le \epsilon_i] \cdot l^{-1} \\
&\quad \left(\mathbb{E}[l(\Lambda) | \epsilon_{i-1} \le \Lambda \le \epsilon_i]\right) \qquad (18) \\
&= l(\lambda_i) \cdot l^{-1}\left(l(\lambda_i)\right) \\
&= \lambda_i \cdot l(\lambda_i),
\end{aligned}
$$

where the inequality step (17) is obtained by applying Jensen's inequality[36,57].

We can now use (13), (14) and (18) to establish an upper bound for the objective function (12):

$$
\mathbb{E}(\Lambda | \text{data}) = \frac{\sum_{i=1}^m \lambda_i' l(\lambda_i') \theta_i}{\sum_{i=1}^m l(\lambda_i) \theta_i} \le \frac{\sum_{i=1}^m \lambda_i l(\lambda_i) \theta_i}{\sum_{i=1}^m l(\lambda_i) \theta_i} \qquad (19)
$$

This upper bound is attained by selecting an $m$-point discrete distribution $f_u(\lambda)$ with probability mass $\theta_i$ at $\lambda = \lambda_i$, for $i = 1, 2, \ldots, m$ (since substituting $f(\cdot)$ from (12) with this $f_u(\cdot)$ yields the rhs result of (19)). As such, maximising this bound reduces to an optimisation problem in the $m$-dimensional space of $(\lambda_1, \lambda_2, \ldots, \lambda_m) \in (\epsilon_0, \epsilon_1] \times (\epsilon_1, \epsilon_2] \times \cdots \times (\epsilon_{m-1}, \epsilon_m]$. This optimisation problem can be solved numerically, yielding a supremum (rather than a maximum) for $\mathcal{S}_\lambda$ in the case when the optimised prior distribution has points located at $\lambda_i = \epsilon_{i-1}$ for $i = 1, 2, \ldots, m$. □

**Proposition 2**. With the notation from Theorem 1, there exist $m$ values $x_1, x_2, \ldots, x_m \in [0, 1]$ such that $\inf \mathcal{S}_\lambda$ is the posterior estimate (4) obtained by using as prior the $(m+1)$-point discrete distribution with probability mass $f(\epsilon_0) = Pr(\lambda = \epsilon_0) = x_1\theta_1$, $f(\epsilon_i) = Pr(\lambda = \epsilon_i) = (1 - x_i)\theta_i + x_{i+1}\theta_{i+1}$ for $1 \le i < m$, and $f(\epsilon_m) = Pr(\lambda = \epsilon_m) = (1 - x_m)\theta_m$.

**Proof**. We reuse the reasoning steps from Proposition 1 up to inequality (17), which we replace with the following alternative inequality derived from the Converse Jensen's Inequality[58,59] and the fact that $g(w)$ is a concave function (cf. Lemma 1):

$$
\begin{aligned}
\lambda_i' l(\lambda_i') &= \mathbb{E}[g(W) | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i] \\
&\ge \frac{l(\epsilon_{i-1}) - \mathbb{E}[W | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i]}{l(\epsilon_{i-1}) - l(\epsilon_i)} g(l(\epsilon_i)) \\
&\quad + \frac{\mathbb{E}[W | \epsilon_{i-1} \le l^{-1}(W) \le \epsilon_i] - l(\epsilon_i)}{l(\epsilon_{i-1}) - l(\epsilon_i)} g(l(\epsilon_{i-1})) \qquad (20) \\
&= \frac{l(\epsilon_{i-1}) - l(\lambda_i)}{l(\epsilon_{i-1}) - l(\epsilon_i)} \epsilon_i l(\epsilon_i) + \frac{l(\lambda_i) - l(\epsilon_i)}{l(\epsilon_{i-1}) - l(\epsilon_i)} \epsilon_{i-1} l(\epsilon_{i-1})
\end{aligned}
$$

We can now establish a lower bound for (12):

$$
\begin{aligned}
\mathbb{E}(\Lambda | \text{data}) &= \frac{\sum_{i=1}^m \lambda_i' l(\lambda_i') \theta_i}{\sum_{i=1}^m l(\lambda_i) \theta_i} \\
&\ge \frac{\sum_{i=1}^m \left( \frac{l(\epsilon_{i-1}) - l(\lambda_i)}{l(\epsilon_{i-1}) - l(\epsilon_i)} \epsilon_i l(\epsilon_i) + \frac{l(\lambda_i) - l(\epsilon_i)}{l(\epsilon_{i-1}) - l(\epsilon_i)} \epsilon_{i-1} l(\epsilon_{i-1}) \right) \theta_i}{\sum_{i=1}^m l(\lambda_i) \theta_i} \qquad (21) \\
&= \frac{\sum_{i=1}^m \left[ \epsilon_i l(\epsilon_i)(1 - x_i)\theta_i + \epsilon_{i-1} l(\epsilon_{i-1}) x_i \theta_i \right]}{\sum_{i=1}^m [l(\epsilon_i)(1 - x_i)\theta_i + l(\epsilon_{i-1}) x_i \theta_i]} \qquad (22)
\end{aligned}
$$

where $x_i$ is defined as:

$$
x_i = \frac{l(\lambda_i) - l(\epsilon_i)}{l(\epsilon_{i-1}) - l(\epsilon_i)} \qquad (23)
$$

The result (22) is essentially in the same form as the result obtained by using a $2m$-point distribution in which, for each interval $[\epsilon_{i-1}, \epsilon_i]$, there are two points located at $\lambda = \epsilon_{i-1}$ and $\lambda = \epsilon_i$ and the probability mass associated with these points is $x_i\theta_i$ and $(1 - x_i)\theta_i$ respectively. Intuitively, $x_i$ is the ratio of splitting the probability mass $\theta_i$ between the two points since, according to (23), $x_i \in [0, 1]$.

Furthermore, the points on the boundaries of two successive intervals are overlapping, which effectively reduces the number of points from $2m$ to $m + 1$. Expanding (22) yields an $(m+1)$-point discrete distribution $f_l(\lambda)$ with probability mass $f_l(\epsilon_0) = x_1\theta_1$, $f_l(\epsilon_i) = (1 - x_i)\theta_i + x_{i+1}\theta_{i+1}$ for $1 \le i < m$ and $f_l(\epsilon_m) = (1 - x_m)\theta_m$. As such, minimising (22) reduces to an $m$-dimensional optimisation problem in $x_1, x_2, \ldots, x_m$, which can be solved numerically given other model parameters. Finally, since (6) requires that $\epsilon_{i-1} < \lambda_i \le \epsilon_i$, we have $0 \le x_i < 1$, and thus the posterior estimate is an infimum (rather than a minimum) of $\mathcal{S}_\lambda$ when the solution of the optimisation problem corresponds to a combination of $x_1, x_2, \ldots, x_m$ values that includes one or more values of 1. □

We can now prove Theorem 1. In the Supplementary Methods 1 of the supplementary material, we use this result to prove Corollaries 1 and 2.

**Proof. Proof of Theorem 1**. Propositions 1 and 2 imply that the set of posterior estimates $\lambda$ over all priors that satisfy the constraints (6) has:

1. the infinum $\lambda_l$ from (7), obtained by using the prior $f(\lambda)$ from Proposition 2 in (4);
2. the supremum $\lambda_u$ from (8), obtained by using the prior $f(\lambda)$ from Proposition 1 in (4). □

**BIPP estimator evaluation**. Figure 5 shows the results of experiments we carried out to evaluate the BIPP estimator in scenarios with $m = 3$ (Fig. 5a–c) and $m = 2$ (Fig. 5d) confidence bounds by varying the characteristics of the partial prior knowledge. For $m = 3$, the upper bound computed by the estimator exhibits a three-stage behaviour as the time over which no singular event occurs increases. These stages correspond to the three $\lambda_u$ regions from (10). They start with a steep $\lambda_u$ decrease for $t < \frac{1}{\epsilon_2}$ in stage 1, followed by a slower $\lambda_u$ decreasing trend for $\frac{1}{\epsilon_2} \le t \le \frac{1}{\epsilon_1}$ in stage 2, and approaching the asymptotic value $\frac{\epsilon_1(\theta_1 + \theta_2)}{\theta_1}$ as the mission progresses through stage 3. Similarly, the lower bound $\lambda_l$ demonstrates a two-stage behaviour, as expected given its two-part definition (9), with the overall value approaching 0 as the mission continues and no singular event modelled by this estimator (e.g., a catastrophic failure) occurs.

Figure 5a shows the behaviour of the estimator for different $\theta_1$ values and fixed $\theta_2$, $\epsilon_1$, and $\epsilon_2$ values. For higher $\theta_1$ values, more probability mass is allocated to the confidence bound $(\epsilon_0, \epsilon_1]$, yielding a steeper decrease in the upper bound $\lambda_u$ and a lower $\lambda_u$ value at the end of the mission. The lower bound $\lambda_l$ presents limited variability across the different $\theta_1$ values, becoming almost constant and close to 0 as $\theta_1$ increases.

A similar decreasing pattern is observed in Fig. 5b, which depicts the results of experiments with $\theta_1$, $\epsilon_1$, and $\epsilon_2$ fixed, and $\theta_2$ variable. The upper bound $\lambda_u$ in the long-term is larger for higher $\theta_2$ values, resulting in a wider posterior estimate bound as $\lambda_u$ converges towards its theoretical asymptotic value.

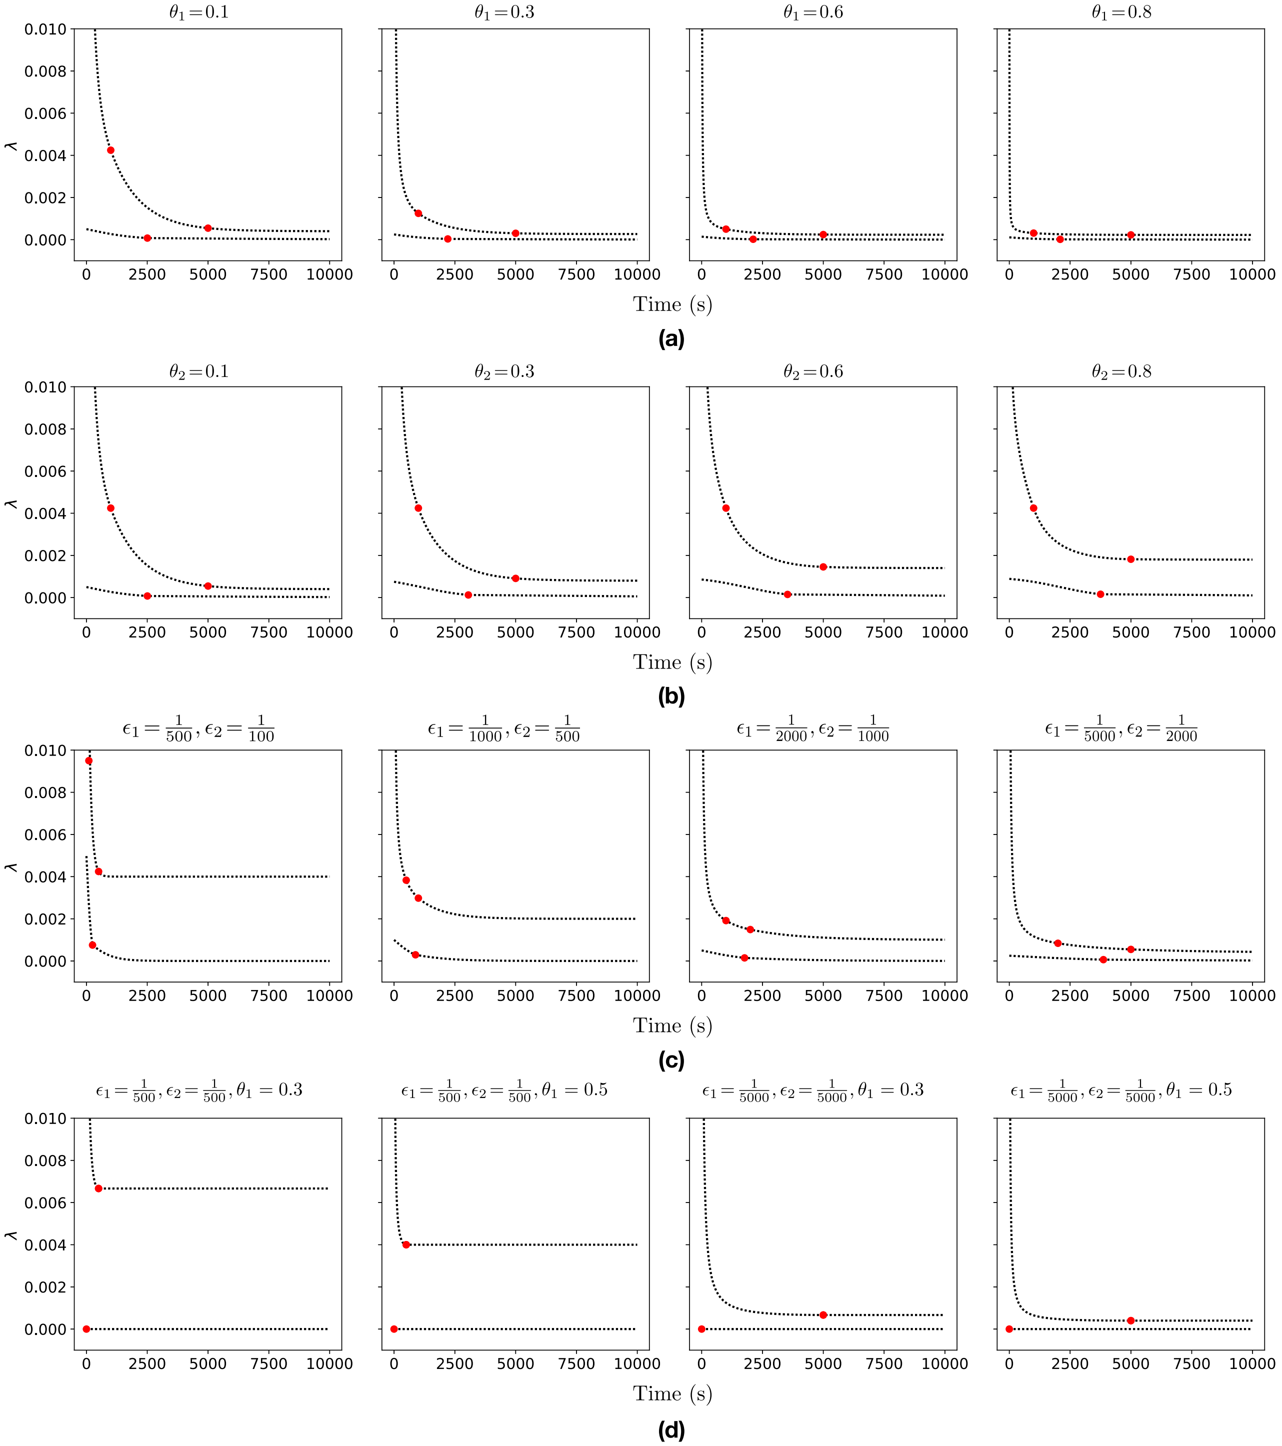

(a)

(b)

(c)

(d)

**Fig. 5 Experimental analysis of the Bayesian inference using partial priors (BIPP) estimator.** Systematic experimental analysis of the BIPP estimator showing the bounds $\lambda_l$ and $\lambda_u$ of the posterior estimates for the occurrence probability of singular events for the duration of a mission. Each plot shows the effect of different partial prior knowledge encoded in (6) on the calculation of the lower (7) and upper (8) posterior estimate bounds. The red circles indicate the time points when the different formulae for the lower and upper bounds in (9) and (10), respectively, become active. **a** BIPP estimator for $m = 3$, $\theta_1 \in \{0.1, 0.3, 0.6, 0.8\}$, $\theta_2 = 0.1$, $\epsilon_1 = \frac{1}{5000}$, $\epsilon_2 = \frac{1}{1000}$. **b** BIPP estimator for $m = 3$, $\theta_1 = 0.1$, $\theta_2 \in \{0.1, 0.3, 0.6, 0.8\}$, $\epsilon_1 = \frac{1}{5000}$, $\epsilon_2 = \frac{1}{1000}$. **c** BIPP estimator for $m = 3$, $\theta_1 = 0.3$, $\theta_2 = 0.3$, $(\epsilon_1, \epsilon_2) \in \left\{\left(\frac{1}{500}, \frac{1}{100}\right), \left(\frac{1}{1000}, \frac{1}{500}\right), \left(\frac{1}{2000}, \frac{1}{1000}\right), \left(\frac{1}{5000}, \frac{1}{2000}\right)\right\}$**d** BIPP estimator for $m = 2$, $\theta_1 \in \{0.3, 0.5\}$, $\theta_2 = 0$, $(\epsilon_1, \epsilon_2) \in \left\{\left(\frac{1}{500}, \frac{1}{500}\right), \left(\frac{1}{5000}, \frac{1}{5000}\right)\right\}$.

Allocating the same probability mass to the confidence bounds, i.e., $\theta_1 = \theta_2 = 0.3$ and changing the prior knowledge bounds $\epsilon_1$ and $\epsilon_2$ affects greatly the behaviour of the BIPP estimator (Fig. 5c). When $\epsilon_1$ and $\epsilon_2$ have relatively high values compared to the duration of the mission (e.g., see the first three plots in Fig. 5c), the upper bound $\lambda_u$ of the BIPP estimator rapidly converges to its asymptotic value, leaving no room for subsequent improvement as the mission progresses. Similarly, the earlier the triggering point for switching between the two parts of the lower bound $\lambda_l$ calculation (9), the earlier $\lambda_l$ reaches a plateau close to 0.

Finally, Fig. 5d shows experimental results for the special scenario comprising only $m = 2$ confidence bounds. In this scenario, replacing $\theta_2 = 0$ in (9) as required by Corollary 2 gives a constant lower bound $\lambda_l = 0$ irrespective of the other BIPP estimator parameters. As expected, the upper bound $\lambda_u$ demonstrates a twofold behaviour, featuring a rapid decrease until $t = \frac{1}{\epsilon 1}$, followed by a steady state behaviour where $\lambda_u = \frac{\epsilon_1}{\theta_1}$.

**Interval Bayesian inference for regular events**. For CTMC transitions that correspond to regular events within the modelled system, we follow the common practice[60] of using a Gamma prior distribution for each uncertain transition rate $\lambda$:

$$f(\lambda) = \Gamma[\lambda; \alpha, \beta] = \frac{\beta^\alpha}{(\alpha-1)!}\lambda^{\alpha-1}e^{-\beta\lambda}. \tag{24}$$

The Gamma distribution is a frequently adopted conjugate prior distribution for the likelihood (2) and, if the prior knowledge assumes an initial value $\lambda^{(0)}$ for the transition rate, the parameters $\alpha > 0$ and $\beta > 0$ must satisfy

$$\mathbb{E}(\Gamma[\lambda; \alpha, \beta]) = \frac{\alpha}{\beta} = \lambda^{(0)}. \tag{25}$$

The posterior value $\lambda^{(t)}$ for the transition rate after observing $n$ transitions within $t$ time units is then obtained by using the prior (24) in the expectation (4), as in the following derivation adapted from classical Bayesian theory[60]:

$$\begin{aligned}
\lambda^{(t)} &= \frac{\int_0^\infty \lambda\left(\frac{(\lambda t)^n}{n!}e^{-\lambda t}\right)\left(\frac{\beta}{(\alpha-1)!}\lambda^{\alpha-1}e^{-\beta\lambda}\right)\mathrm{d}\lambda}{\int_0^\infty \left(\frac{(\lambda t)^n}{n!}e^{-\lambda t}\right)\left(\frac{\beta}{(\alpha-1)!}\lambda^{\alpha-1}e^{-\beta\lambda}\right)\mathrm{d}\lambda}\\
&= \frac{\int_0^\infty \lambda^{n+\alpha}e^{-\lambda(t+\beta)}\mathrm{d}\lambda}{\int_0^\infty \lambda^{n+\alpha-1}e^{-\lambda(t+\beta)}\mathrm{d}\lambda} = \frac{\int_0^\infty \lambda^{n+\alpha}\left(\frac{e^{-\lambda(t+\beta)}}{-(t+\beta)}\right)'\mathrm{d}\lambda}{\int_0^\infty \lambda^{n+\alpha-1}e^{-\lambda(t+\beta)}\mathrm{d}\lambda}\\
&= \frac{\left(\lambda^{n+\alpha}\frac{e^{-\lambda(t+\beta)}}{-(t+\beta)}\right)\Big|_0^\infty - \int_0^\infty (n+\alpha)\lambda^{n+\alpha-1}\frac{e^{-\lambda(t+\beta)}}{-(t+\beta)}\mathrm{d}\lambda}{\int_0^\infty \lambda^{n+\alpha-1}e^{-\lambda(t+\beta)}\mathrm{d}\lambda}\\
&= \frac{0 + \frac{n+\alpha}{t+\beta}\int_0^\infty \lambda^{n+\alpha-1}e^{-\lambda(t+\beta)}\mathrm{d}\lambda}{\int_0^\infty \lambda^{n+\alpha-1}e^{-\lambda(t+\beta)}\mathrm{d}\lambda} = \frac{n+\alpha}{t+\beta} = \frac{n+\beta\lambda^{(0)}}{t+\beta}\\
&= \frac{\beta}{t+\beta}\lambda^{(0)} + \frac{t}{t+\beta}\frac{n}{t} = \frac{t^{(0)}}{t+t^{(0)}}\lambda^{(0)} + \frac{t}{t+t^{(0)}}\frac{n}{t},
\end{aligned} \tag{26}$$

where $t^{(0)} = \beta$. This notation reflects the way in which the posterior rate $\lambda^{(t)}$ is computed as a weighted sum of the mean rate $\frac{n}{t}$ observed over a time period $t$, and of the prior $\lambda^{(0)}$ deemed as trustworthy as a mean rate calculated from observations over a time period $t^{(0)}$. When $t^{(0)} \ll t$ (either because we have low trust in the prior $\lambda^{(0)}$ and thus $t^{(0)} \simeq 0$, or because the system was observed for a time period $t$ that is much longer than $t^{(0)}$), the posterior (26) reduces to the maximum likelihood estimator, i.e. $\lambda^{(t)} \simeq \frac{n}{t}$. In this scenario, the observations fully dominate the estimator (26), with no contribution from the prior.

The selection of suitable values for the parameters $t^{(0)}$ and $\lambda^{(0)}$ of the traditional Bayesian estimator (26) is very challenging. What constitutes a suitable choice often depends on unknown attributes of the environment, or several domain experts may each propose different values for these parameters. In line with recent advances in imprecise probabilistic modelling[34,35,61], we address this challenge by defining a robust transition rate estimator for Bayesian inference using imprecise probability with sets of priors (IPSP). The IPSP estimator uses ranges $[\underline{t}^{(0)}, \overline{t}^{(0)}]$ and $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$ (corresponding to the environmental uncertainty, or to input obtained from multiple domain experts) for the two parameters instead of point values.

The following theorem quantifies the uncertainty that the use of parameter ranges for $t^{(0)}$ and $\lambda^{(0)}$ induces on the posterior rate (26). This theorem specialises and builds on generalised Bayesian inference results[34] that we adapt for the estimation of CTMC transition rates.

**Theorem 2**. Given uncertain prior parameters $t^{(0)} \in [\underline{t}^{(0)}, \overline{t}^{(0)}]$ and $\lambda^{(0)} \in [\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$, the posterior rate $\lambda^{(t)}$ from (26) can range in the interval $[\underline{\lambda}^{(t)}, \overline{\lambda}^{(t)}]$, where:

$$\underline{\lambda}^{(t)} = \begin{cases} \frac{\overline{t}^{(0)}\underline{\lambda}^{(0)}+n}{\overline{t}^{(0)}+t}, & \textit{if } \frac{n}{t} \geq \underline{\lambda}^{(0)} \\ \frac{\underline{t}^{(0)}\underline{\lambda}^{(0)}+n}{\underline{t}^{(0)}+t}, & \text{otherwise} \end{cases} \tag{27}$$

and

$$\overline{\lambda}^{(t)} = \begin{cases} \frac{\overline{t}^{(0)}\overline{\lambda}^{(0)}+n}{\overline{t}^{(0)}+t}, & \textit{if } \frac{n}{t} \leq \overline{\lambda}^{(0)} \\ \frac{\underline{t}^{(0)}\overline{\lambda}^{(0)}+n}{\underline{t}^{(0)}+t}, & \text{otherwise} \end{cases}. \tag{28}$$

**Proof**. To find the extrema for the posterior rate $\lambda^{(t)}$, we first differentiate (26) with respect to $\lambda^{(0)}$:

$$\frac{d}{d\lambda^{(0)}}\left(\lambda^{(t)}\right) = \frac{t^{(0)}}{t+t^{(0)}}.$$

As $t^{(0)} > 0$ and $t > 0$, this derivative is always positive, so

$$\underline{\lambda}^{(t)} = \min_{t^{(0)} \in [\underline{t}^{(0)}, \overline{t}^{(0)}]} \frac{t^{(0)}\underline{\lambda}^{(0)} + n}{t^{(0)} + t} \tag{29}$$

and

$$\overline{\lambda}^{(t)} = \max_{t^{(0)} \in [\underline{t}^{(0)}, \overline{t}^{(0)}]} \frac{t^{(0)}\overline{\lambda}^{(0)} + n}{t^{(0)} + t}. \tag{30}$$

We now differentiate the quantity that needs to be minimised in (29) with respect to $t^{(0)}$:

$$\begin{aligned}
\frac{d}{dt^{(0)}}\left(\frac{t^{(0)}\underline{\lambda}^{(0)} + n}{t^{(0)} + t}\right) &= \frac{\underline{\lambda}^{(0)}(t^{(0)} + t) - (t^{(0)}\underline{\lambda}^{(0)} + n)\cdot 1}{(t^{(0)} + t)^2}\\
&= \frac{\underline{\lambda}^{(0)}t - n}{(t^{(0)} + t)^2},
\end{aligned}$$

As this derivative is non-positive for $\underline{\lambda}^{(0)} \in \left(0, \frac{n}{t}\right]$ and positive for $\underline{\lambda}^{(0)} > \frac{n}{t}$, the minimum from (29) is attained for $t^0 = \overline{t}^{(0)}$ in the former case, and for $t^0 = \underline{t}^{(0)}$ in the latter case, which yields the result from (27). Similarly, the derivative of the quantity to maximise in (30), i.e.,

$$\frac{d}{dt^{(0)}}\left(\frac{t^{(0)}\overline{\lambda}^{(0)} + n}{t^{(0)} + t}\right) = \frac{\overline{\lambda}^{(0)}t - n}{(t^{(0)} + t)^2},$$

is non-positive for $\overline{\lambda}^{(0)} \in \left(0, \frac{n}{t}\right]$ and positive for $\overline{\lambda}^{(0)} > \frac{n}{t}$, so the maximum from (30) is attained for $t^0 = \underline{t}^{(0)}$ in the former case, and for $t^0 = \overline{t}^{(0)}$ in the latter case, which yields the result from (28) and completes the proof. □

**IPSP estimator evaluation**. Figure 6 shows the results of experiments we performed to analyse the behaviour of the IPSP estimator in scenarios with varying ranges for the prior knowledge $[\underline{t}^{(0)}, \overline{t}^{(0)}]$ and $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$. A general observation is that the posterior rate intervals $[\underline{\lambda}^{(t)}, \overline{\lambda}^{(t)}]$ become narrower as the mission progresses, irrespective of the level of trust assigned to the prior knowledge, i.e., across all columns of plots (which correspond to different $[\underline{t}^{(0)}, \overline{t}^{(0)}]$ intervals) from Fig. 6a. Nevertheless, this trust

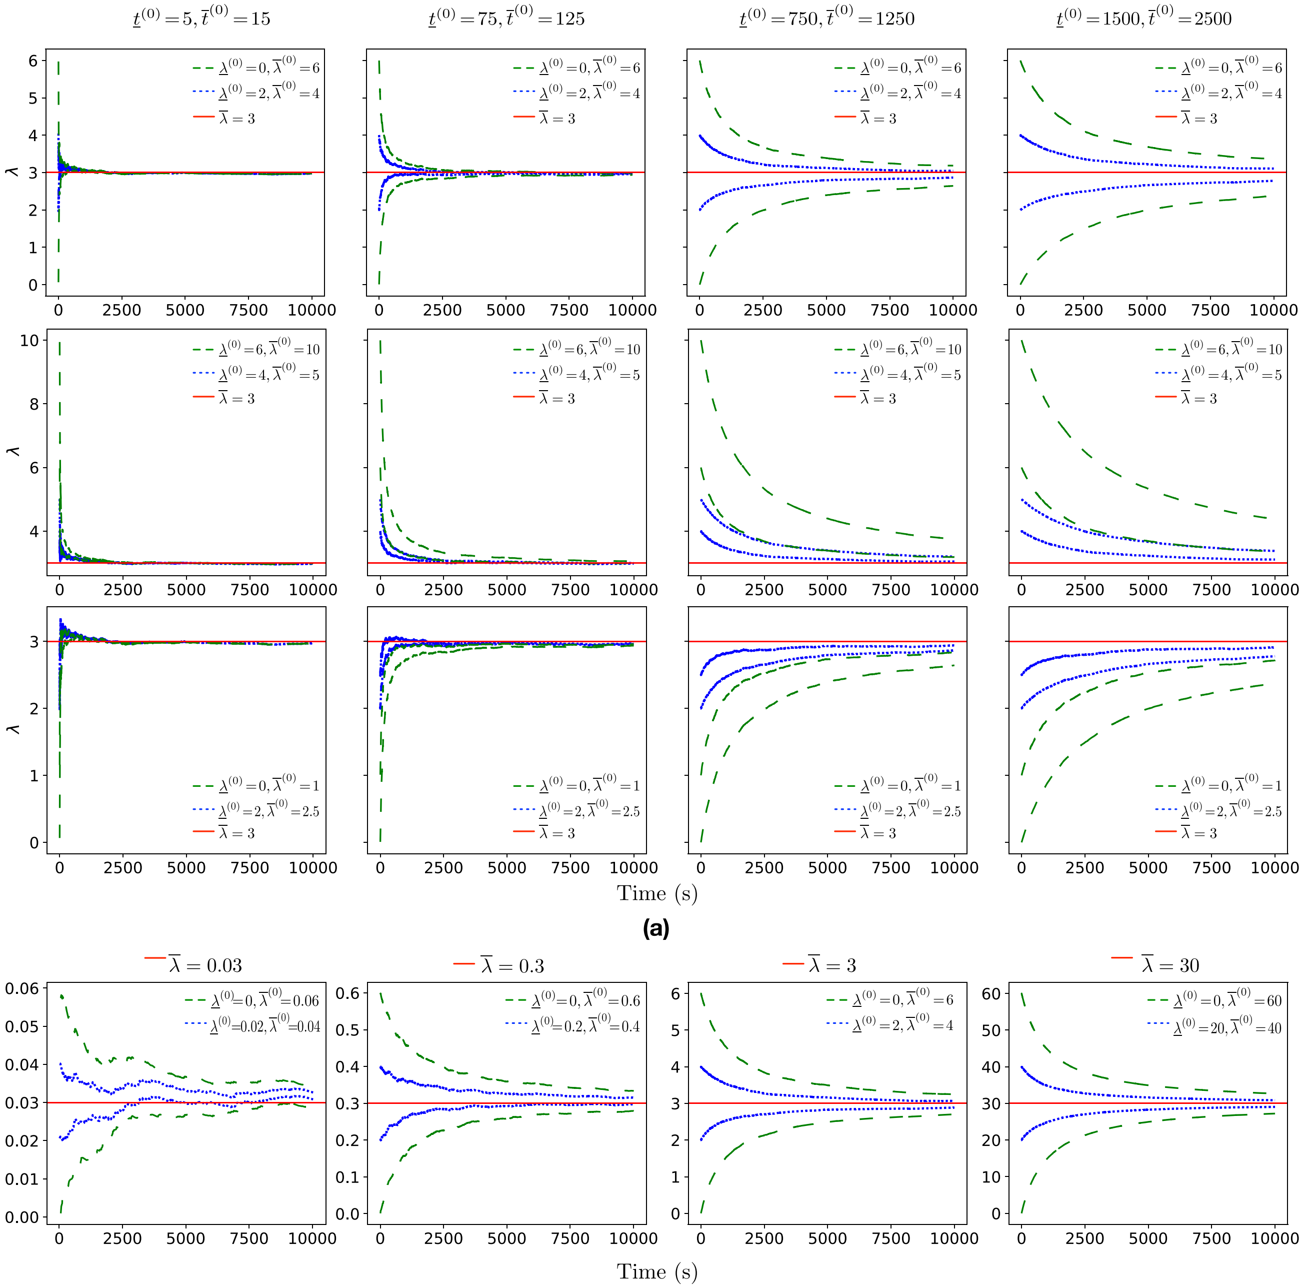

(a)

(b)

**Fig. 6 Experimental analysis of the Bayesian inference using imprecise probability with sets of priors (IPSP) estimator showing the bounded posterior estimation for regular events. a** IPSP estimator results showing the impact of different sets of priors $[\underline{t}^{(0)}, \overline{t}^{(0)}]$ and $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$. In each plot, the blue dotted line ( ···· ) and green dashed line ( — — — ) show the posterior estimation bounds $\underline{\lambda}^{(t)}$ and $\overline{\lambda}^{(t)}$ for narrow and wide $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$ intervals, respectively. Each column of plots corresponds to assigning different strength to the prior knowledge, ranging from uninformative (leftmost column) to strong belief (rightmost column). The first row shows scenarios in which the actual rate $\bar{\lambda} = 3$ belongs to the prior knowledge interval $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$. In the second and third rows, the prior intervals overestimate and underestimate $\bar{\lambda}$, respectively. **b** IPSP estimator results illustrating the behaviour of IPSP across different actual rate values $\bar{\lambda} \in \{0.03, 0.3, 3, 30\}$. The experiments were carried out for $[\underline{t}^{(0)}, \overline{t}^{(0)}] = [1000, 1000]$ and included both narrow and wide $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$ intervals, which are shown in blue dotted lines ( ···· ) and green dashed lines ( — — — ), respectively. In all experiments, the unknown actual rate $\bar{\lambda}$ was in the prior interval $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$.

level affects how the estimator incorporates observations into the calculation of the posterior interval. When the trust in the prior knowledge is weak (in the plots from the leftmost columns of Fig. 6a), the impact of the prior knowledge on the posterior estimation is low, and the IPSP calculation is heavily influenced by the observations, resulting in a narrow interval. In contrast, when the trust in the prior knowledge is stronger (in the plots from the rightmost columns), the contribution of the prior knowledge to the posterior estimation becomes higher, and the IPSP estimator produces a wider interval.

In the experiments from the first row of plots in Fig. 6a, the (unknown) actual rate $\bar{\lambda} = 3$ belongs to the prior knowledge interval $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$. As a result, the posterior rate interval $[\underline{\lambda}^{(t)}, \overline{\lambda}^{(t)}]$ progressively becomes narrower, approximating $\overline{\lambda}$ with high accuracy. As expected, the narrower prior knowledge (blue dotted line) produces a narrower posterior rate interval than the wider and more conservative prior knowledge (green dashed line).

When the prior knowledge interval $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$ overestimates or underestimates the actual rate $\overline{\lambda}$ (second and third rows of plots from Fig. 6a, respectively), the ability of IPSP to adapt its estimations to reflect the observations heavily depends on the characteristics of the sets of priors. For example, if the width of the prior knowledge $[\underline{\lambda}^{(0)}, \overline{\lambda}^{(0)}]$ is close to $\overline{\lambda}$ and $t^{(0)} \ll t$, then IPSP more easily approaches $\overline{\lambda}$, as shown by the narrow prior knowledge (blue dotted line) in Fig. 6a for $[\underline{t}^{(0)}, \overline{t}^{(0)}] \in \{[5, 15], [75, 125], [750, 1250]\}$. In contrast, wider narrow prior knowledge (green dashed line) combined with higher levels of trust in the prior, e.g., $[\underline{t}^{(0)}, \overline{t}^{(0)}] \in \{[1500, 2500]\}$, entails that more observations are needed for the posterior rate to approach the actual rate $\overline{\lambda}$. When the actual rate is, in addition, nonstationary, change-point detection methods can be employed to identify these changes[62,63] and recalibrate the IPSP estimator. Finally, Fig. 6b shows the behaviour of IPSP for different actual rate $\overline{\lambda}$ values, i.e., $\overline{\lambda} \in \{0.03, 0.3, 3, 30\}$. As $\overline{\lambda}$ increases, more observations are produced in the same time period, resulting in a smoother and narrower posterior bound estimate.

## Data availability

The data supporting the RBV findings and a video of the robotic mission in simulation are available at https://gerasimou.github.io/RBV.

## Code availability

All code developed in this project is freely available at http://github.com/gerasimou/RBV.

Received: 19 April 2023; Accepted: 11 January 2024;
Published online: 27 January 2024

## References

1. The Headquarters for Japan's Economic Revitalization. New Robot Strategy: Japan's Robot Strategy. *Prime Minister's Office of Japan* (2015).
2. SPARC–The Partnership for Robotics in Europe. Robotics 2020 multi-annual roadmap for robotics in Europe. *eu-robotics* (2016).
3. Science and Technology Committee. Robotics and Artificial Intelligence. *Committee Reports of UK House of Commons* (2016).
4. Christensen, H. et al. A roadmap for us robotics–from internet to robotics 2020 edition. *Found. Trends Robot.* **8**, 307–424 (2021).
5. Richardson, R. et al. Robotic and autonomous systems for resilient infrastructure. *UK-RAS White Papers© UK-RAS* (2017).
6. UK Robotics & Autonomous Systems Network. Space Robotics & Autonomous Systems: Widening the horizon of space exploration. *UK-RAS White Papers© UK-RAS* (2018).
7. Lane, D., Bisset, D., Buckingham, R., Pegman, G. & Prescott, T. New foresight review on robotics and autonomous systems. Tech. Rep. No. 2016.1. (Lloyd's Register Foundation, London, UK, 2016).
8. Calinescu, R. et al. Engineering trustworthy self-adaptive software with dynamic assurance cases. *IEEE Trans. Softw. Eng.* **44**, 1039–1069 (2017).
9. Robu, V., Flynn, D. & Lane, D. Train robots to self-certify as safe. *Nature* **553**, 281–281 (2018).
10. Calinescu, R., Ghezzi, C., Kwiatkowska, M. & Mirandola, R. Self-adaptive software needs quantitative verification at runtime. *Commun. ACM* **55**, 69–77 (2012).
11. International Nuclear Safety Advisory Group. Defence in Depth in Nuclear Safety (INSAG 10) (1996).
12. Kwiatkowska, M. Quantitative verification: models, techniques and tools. In *Proc. 6th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*, 449–458 (ACM Press, 2007).
13. Katoen, J.-P. The Probabilistic Model Checking Landscape. In *Proc. of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '16, 31–45 (ACM, New York, NY, USA, 2016). https://doi.org/10.1145/2933575.2934574.
14. Legay, A., Delahaye, B. & Bensalem, S. Statistical model checking: an overview. In (eds Barringer, H. et al.) *Runtime Verification*, vol. 6418 of *LNCS*, 122–135 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2010).
15. Calinescu, R. et al. Formal verification with confidence intervals to establish quality of service properties of software systems. *IEEE Trans. Reliab.* **65**, 107–125 (2016).
16. Kwiatkowska, M., Norman, G. & Parker, D. Probabilistic model checking and autonomy. *Annu. Rev. Control Robot. Autonomous Syst.* **5**, 385–410 (2022).
17. Lacerda, B., Faruq, F., Parker, D. & Hawes, N. Probabilistic planning with formal performance guarantees for mobile service robots. *Int. J. Robot. Res.* **38**, 1098–1123 (2019).
18. Nardone, V., Santone, A., Tipaldi, M. & Glielmo, L. Probabilistic model checking applied to autonomous spacecraft reconfiguration. In *IEEE Metrology for Aerospace (MetroAeroSpace)*, 556–560 (IEEE, 2016).
19. Fraser, D. et al. Collaborative models for autonomous systems controller synthesis. *Form. Asp. Comput.* **32**, 157–186 (2020).
20. Liu, W. & Winfield, A. F. Modeling and optimization of adaptive foraging in swarm robotic systems. *Int. J. Robot. Res.* **29**, 1743–1760 (2010).
21. Brim, L., Ceska, M., Drazan, S. & Safranek, D. Exploring parameter space of stochastic biochemical systems using quantitative model checking. In *Computer Aided Verification (CAV)*, 107–123 (2013).
22. Ceska, M., Pilar, P., Paoletti, N., Brim, L. & Kwiatkowska, M. PRISM-PSY: Precise GPU-accelerated parameter synthesis for stochastic systems. In (eds Chechik, M. & Raskin, J.-F.) *Tools and Algorithms for the Construction and Analysis of Systems*, vol. 9636 of *LNCS*, 367–384 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2016).
23. Avizienis, A., Laprie, J., Randell, B. & Landwehr, C. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Secur. Comput.* **1**, 11–33 (2004).
24. Lane, D. M. et al. PANDORA-persistent autonomy through learning, adaptation, observation and replanning. *IFAC-PapersOnLine* **48**, 238–243 (2015).
25. Myhr, A., Bjerkseter, C., Ågotnes, A. & Nygaard, T. A. Levelised cost of energy for offshore floating wind turbines in a life cycle perspective. *Renew. Energy* **66**, 714–728 (2014).
26. Benjamin, M. R., Schmidt, H., Newman, P. M. & Leonard, J. J. Autonomy for unmanned marine vehicles with MOOS-IvP. In (ed Seto, M. L.) *Marine Robot Autonomy*, 47–90 (Springer, 2013). https://doi.org/10.1007/978-1-4614-5659-9_2.
27. Epifani, I., Ghezzi, C., Mirandola, R. & Tamburrelli, G. Model evolution by run-time parameter adaptation. In *Proc. of the 31st Int. Conf. on Software Engineering*, 111–121 (IEEE, Washington, DC, USA, 2009).
28. Filieri, A., Ghezzi, C. & Tamburrelli, G. A formal approach to adaptive software: continuous assurance of non-functional requirements. *Form. Asp. Comput.* **24**, 163–186 (2012).
29. Calinescu, R., Rafiq, Y., Johnson, K. & Bakir, M. E. Adaptive model learning for continual verification of non-functional properties. In *Proc. of the 5th Int. Conf. on Performance Engineering*, 87–98 (ACM, NY, USA, 2014).
30. Filieri, A., Grunske, L. & Leva, A. Lightweight adaptive filtering for efficient learning and updating of probabilistic models. In *Proc. of the 37th Int. Conf. on Software Engineering*, 200–211 (IEEE Press, Piscataway, NJ, USA, 2015).
31. Calinescu, R., Johnson, K. & Paterson, C. FACT: A probabilistic model checker for formal verification with confidence intervals. In (eds Chechik, M. & Raskin, J.-F.) *Tools and Algorithms for the Construction and Analysis of Systems*, 540–546 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2016).
32. Calinescu, R., Češka, M., Gerasimou, S., Kwiatkowska, M. & Paoletti, N. RODES: a robust-design synthesis tool for probabilistic systems. In *Quantitative Evaluation of Systems: 14th International Conference, QEST 2017, Berlin, Germany, September 5-7, 2017, Proceedings 14*, 304–308 (Springer, 2017).
33. Zhao, X. et al. Probabilistic model checking of robots deployed in extreme environments. In *Proc. of the 33rd AAAI Conference on Artificial Intelligence*, vol. 33, 8076–8084 (Honolulu, Hawaii, USA, 2019).
34. Walter, G. & Augustin, T. Imprecision and prior-data conflict in generalized Bayesian inference. *J. Stat. Theory Pract.* **3**, 255–271 (2009).
35. Walter, G., Aslett, L. & Coolen, F. P. A. Bayesian nonparametric system reliability using sets of priors. *Int. J. Approx. Reason.* **80**, 67–88 (2017).
36. Bishop, P., Bloomfield, R., Littlewood, B., Povyakalo, A. & Wright, D. Toward a formalism for conservative claims about the dependability of software-based systems. *IEEE Trans. Softw. Eng.* **37**, 708–717 (2011).
37. Strigini, L. & Povyakalo, A. Software fault-freeness and reliability predictions. In (eds Bitsch, F., Guiochet, J. & Kaâniche, M.) *Computer Safety, Reliability, and Security*, vol. 8153 of *LNCS*, 106–117 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2013).
38. Zhao, X., Salako, K., Strigini, L., Robu, V. & Flynn, D. Assessing safety-critical systems from operational testing: a study on autonomous vehicles. *Inf. Softw. Technol.* **128**, 106393 (2020).

39. Ishikawa, A. et al. The max-min delphi method and fuzzy delphi method via fuzzy integration. *Fuzzy Sets Syst.* **55**, 241–253 (1993).
40. Flyvbjerg, B. Curbing optimism bias and strategic misrepresentation in planning: reference class forecasting in practice. *Eur. Plan. Stud.* **16**, 3–21 (2008).
41. Araujo, H., Mousavi, M. R. & Varshosaz, M. Testing, validation, and verification of robotic and autonomous systems: a systematic review. *ACM Trans. Softw. Eng. Methodol.* **32**, 1–61 (2023).
42. Luckcuck, M., Farrell, M., Dennis, L. A., Dixon, C. & Fisher, M. Formal specification and verification of autonomous robotic systems: a survey. *ACM Comput. Surv.* **52**, 1–41 (2019).
43. Gleirscher, M., Foster, S. & Woodcock, J. New opportunities for integrated formal methods. *ACM Comput. Surv.* **52**, 1–36 (2019).
44. Gerasimou, S., Calinescu, R., Shevtsov, S. & Weyns, D. UNDERSEA: an exemplar for engineering self-adaptive unmanned underwater vehicles. In *IEEE/ACM 12th Int. Symp. on Software Engineering for Adaptive and Self-Managing Systems*, 83–89 (2017).
45. Younes, H. L., Kwiatkowska, M., Norman, G. & Parker, D. Numerical vs. statistical probabilistic model checking. *Int. J. Softw. Tools Technol. Transf.* **8**, 216–228 (2006).
46. Zhang, L., Hermanns, H. & Jansen, D. N. Logic and model checking for hidden Markov models. In (ed Wang, F.) *Formal Techniques for Networked and Distributed Systems - FORTE 2005*, 98–112 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2005).
47. Hernández, N., Eder, K., Magid, E., Savage, J. & Rosenblueth, D. A. Marimba: a tool for verifying properties of hidden markov models. In (eds Finkbeiner, B., Pu, G. & Zhang, L.) *Automated Technology for Verification and Analysis*, 201–206 (Springer International Publishing, Cham, 2015).
48. Wei, W., Wang, B. & Towsley, D. Continuous-time hidden Markov models for network performance evaluation. *Perform. Eval.* **49**, 129–146 (2002). Performance 2002.
49. Baier, C., Haverkort, B., Hermanns, H. & Katoen, J. P. Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Softw. Eng.* **29**, 524–541 (2003).
50. Kwiatkowska, M., Norman, G. & Parker, D. Stochastic model checking. In *International Conference on Formal Methods for Performance Eval*uation. 220–270 (2007).
51. Aziz, A., Sanwal, K., Singhal, V. & Brayton, R. Verifying continuous time Markov chains. In *Computer Aided Verification*, 269–276 (Springer, 1996).
52. Kwiatkowska, M., Norman, G. & Parker, D. PRISM 4.0: Verification of probabilistic real-time systems. In *Proc. of the 23rd Int. Conf. on Computer Aided Verification*, vol. 6806 of *LNCS*, 585–591 (Springer, 2011).
53. Dehnert, C., Junges, S., Katoen, J.-P. & Volk, M. A Storm is coming: a modern probabilistic model checker. In *29th International Conference on Computer Aided Verification (CAV)*, 592–600 (2017).
54. Calinescu, R., Ceska, M., Gerasimou, S., Kwiatkowska, M. & Paoletti, N. Efficient synthesis of robust models for stochastic systems. *J. Syst. Softw.* **143**, 140–158 (2018).
55. International Electrotechnical Commission. IEC 61508—Functional safety of electrical/electronic/programmable electronic safety-related systems (2010).
56. Gradshteyn, I. S. & Ryzhik, I. M. Definite integrals of elementary functions. In (eds Zwillinger, D. & Moll, V.) *Table of Integrals, Series, and Products* (Elsevier Science, 2015), 8th edn.
57. Jensen, J. L. W. V. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math.* **30**, 175–193 (1906).
58. Lah, P. & Ribarič, M. Converse of Jensen's inequality for convex functions. *Publ. Elektroteh. Fak. Ser. Mat. Fiz.* **412/460**, 201–205 (1973).
59. KlaričIć Bakula, M., Pečarić, J. & Perić, J. On the converse Jensen inequality. *Appl. Math. Comput.* **218**, 6566–6575 (2012).
60. Bernardo, J. M. & Smith, A. F. M. *Bayesian Theory* (Wiley, 1994).
61. Krpelik, D., Coolen, F. P. & Aslett, L. J. Imprecise probability inference on masked multicomponent system. In *International Conference Series on Soft Methods in Probability and Statistics*, 133–140 (Springer, 2018).
62. Epifani, I., Ghezzi, C. & Tamburrelli, G. Change-point detection for black-box services. In *Proc. of the 18th ACM SIGSOFT Int. Symp. on Foundations of Software Engineering*, FSE '10, 227–236 (ACM, New York, NY, USA, 2010).
63. Zhao, X., Calinescu, R., Gerasimou, S., Robu, V. & Flynn, D. Interval change-point detection for runtime probabilistic model checking. In *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 163–174 (IEEE, 2020).

## Acknowledgements

This project has received funding from the ORCA-Hub PRF project 'COVE', the Assuring Autonomy International Programme, the UKRI project EP/V026747/1 'Trustworthy Autonomous Systems Node in Resilience', and the European Union's Horizon 2020 project SESAME (grant agreement No. 101017258).

## Author contributions

X.Z.: Conceptualisation, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Software, Supervision, Validation, Writing—original draft, Writing—review and editing; S.G.: Conceptualisation, Data curation, Funding acquisition, Investigation, Methodology, Project administration, Software, Supervision, Validation, Visualisation, Writing—original draft, Writing—review and editing; R.C.: Conceptualisation, Formal analysis, Funding acquisition, Methodology, Project administration, Supervision, Visualisation, Writing—original draft, Writing—review and editing; C.I.: Data curation, Investigation, Validation, Visualisation, Writing—original draft, Writing—review and editing; V.R.: Conceptualisation, Funding acquisition, Methodology, Project administration, Supervision, Writing—review and editing; D.F.: Conceptualisation, Funding acquisition, Methodology, Project administration, Supervision, Writing—review and editing.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at https://doi.org/10.1038/s44172-024-00162-y.

**Correspondence** and requests for materials should be addressed to Xingyu Zhao or Simos Gerasimou.

**Peer review information** *Communications Engineering* thanks Koorosh Aslansefat and the other, anonymous, reviewer for their contribution to the peer review of this work. Primary Handling Editors: Alessandro Rizzo and Mengying Su.

**Reprints and permission information** is available at http://www.nature.com/reprints

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.