Updating Autonomous Underwater Vehicle Risk Based on the Effectiveness of Failure Prevention and Correction

MARIO P. BRITO

University of Southampton, Southampton, Hampshire, United Kingdom

GWYN GRIFFITHS

Autonomous Analytics, Southampton, Hampshire, United Kingdom

(Manuscript received 27 December 2016, in final form 24 January 2018)

ABSTRACT

Autonomous underwater vehicles (AUVs) have proven to be feasible platforms for marine observations. Risk and reliability studies on the performance of these vehicles by different groups show a significant difference in reliability, with the observation that the outcomes depend on whether the vehicles are operated by developers or nondevelopers. This paper shows that this difference in reliability is due to the failure prevention and correction procedures—risk mitigation—put in place by developers. However, no formalization has been developed for updating the risk profile based on the expected effectiveness of the failure prevention and correction process. A generic Bayesian approach for updating the risk profile is presented, based on the probability of failure prevention and correction and the number of subsequent deployments on which the failure does not occur. The approach, which applies whether the risk profile is captured in a parametric or nonparametric survival model, is applied to a real case study of the International Submarine Engineering Ltd. (ISE) Explorer AUV.

1. Introduction

Autonomous underwater vehicles (AUVs) are mechatronic robotic systems that are able to navigate underwater while untethered from any other system. For the purpose of this study, underwater gliders, which use a buoyancy change engine for propulsion but otherwise share many of the attributes of AUVs, are considered AUVs. With ships being expensive to operate and satellites unable to observe many ocean features, or restricted to observing the near surface, AUVs are an effective technology to sample the ocean (Singh et al. 2004; Webb et al. 2001; Eriksen et al. 2001; Rudnick et al. 2016).

Early studies on propeller-driven AUV risk and reliability presented analysis methodologies with examples from a range of deployments (Podder et al. 2004; Griffiths and Trembanis 2007; Griffiths et al. 2009; Brito et al. 2010, 2012). More recently, two independent studies of underwater glider reliability have shown a significant difference in performance between gliders maintained and deployed by their developers and those deployed by purchasers (Brito et al. 2014). Rudnick et al. (2016) examined the operation of the Spray underwater glider by the development and operations team at the Scripps Institution of Oceanography. Their survival analysis concluded that for a 100-day mission, the probability of survival for the Spray glider was 0.83. For this calculation the authors considered the faults that led to premature mission abort albeit in some cases the mission was not aborted, because the main aim was to demonstrate a target mission length rather than to gather scientific data. In contrast, an analysis of commercially available gliders, operated by nondevelopers, concluded that the probability of a deep glider surviving a 90-day mission without premature mission abort was 0.5 (Brito et al. 2014). Differences in survival estimates have also been observed for the risk of vehicle loss, with Rudnick et al. (2016) reporting a survival of 0.95 for a 100-day mission, and Brito et al. (2014) reporting a survival of 0.8 for a 100-day mission. In their survival analysis, with respect to vehicle loss, Rudnick et al. (2016) considered faults that led to a loss of control over buoyancy and vehicle loss as failures. In the study by Brito et al. (2014),

Denotes content that is immediately available upon publication as open access.

Corresponding author: Mario P. Brito, m.p.brito@soton.ac.uk

DOI: 10.1175/JTECH-D-16-0252.1

the authors considered vehicle loss as failures. Both studies argue that understanding and eliminating failure modes are key to increasing the probability of successful mission completion and survival.

At this stage it is important to distinguish between *failure* and *faults*, as these two terms are used in the manuscript. Our definition of *failure* is aligned with that adopted by the British Standard (BSI 1991, p. 5), which states that *failure* is ''the termination of the ability of an item to perform a required function.'' A failure is a result of a component fault or human error. A component fault is caused by a defect and human error is caused by a person's lapse, slip, mistake, or violation (Reason 1990). For simplicity, in this paper we use *fault* to encapsulate a component defect and a human error. In the work presented by Brito et al. (2014) and Rudnick et al. (2016), mission abort and AUV loss were considered mission failures. In this paper *failure* is defined as the termination of the ability of an AUV to perform the required function that can potentially lead to AUV loss. Generally, *failure mitigation* (risk mitigation) is defined as the process of annulling the consequence of failure or its likelihood of occurrence (Subramanian et al. 1996). In this paper *failure mitigation* (risk mitigation) is achieved by reducing the likelihood of failure occurrence.

While both Brito et al. (2014) and Rudnick et al. (2016) give emphasis to the role of mitigation through failure prevention and correction, neither presents an analytical framework for updating the risk profile based on a structured assessment of the understanding and elimination of failure modes and the subsequent effect on field results. Such an analytical approach would have to be based on probability theory, as is proposed in this paper.

The novelty of the technology makes it impossible to obtain past data on the probability of failure mitigation. Therefore, the assessment of the probability of failure mitigation for this type of technology must rely on expert subjective judgment. The same approach has been adopted in space exploration (Feather and Cornford 2003).

Feather and Cornford (2003) presented a hazard management framework to monitor and update the likelihood of the occurrence of design failure modes occurring. The system, Defect Detection and Prevention (DDP), is a probabilistic model. The key assumption is that each failure may have a number of prevention, analysis, control, and test (PACT) methods. The efficiency of each PACT in mitigating the failure is assessed by a group of experts in the field. The DDP system considers that multiple PACTs may have adverse or positive effects on a failure mode. PACTs are not independent, and they may introduce a failure into the system. The probability model aggregates all these effects to quantify the likelihood of failure mitigation. Brito et al. (2012) also used probability for modeling failure mitigation. However, in contrast to Feather and Cornford (2003), the mitigation actions were considered to be independent.

In this paper we present a Bayesian approach to updating the risk profile of an AUV, based on the preimplementation perceived effectiveness of the mitigation by subject matter experts and the observed performance during subsequent missions. We present a case study of the International Submarine Engineering Ltd. (ISE) Explorer AUV to illustrate the application of the method, based on the initial failure and survivability data presented in Brito et al. (2012).

This paper is organized as follows. Section 2 presents a summary of the methods used for estimating AUV survival. Section 3 presents the data of the ISE Explorer campaigns in the Arctic in 2010 and 2011. Section 4 presents the method proposed for updating the risk profile of an autonomous vehicle based on the probability of failure mitigation and field results. Section 5 presents the application of the method to the ISE case study. Section 6 presents the conclusions.

2. Survival analysis for autonomous underwater vehicles

For successful AUV missions, the observation that we gather is that the vehicle survived at least time t or distance $x(t)$; this is the total mission time or mission length. In statistical survival modeling, this observation where the end, or the last reading, did not result in death is denoted as censored. For AUVs, successful missions are modeled as right censored, which means that the vehicle has survived at least the time or distance traveled by the vehicle. Some missions, however, end in failure. In statistical modeling, this event is represented as death. A survival function or distribution is a mathematical function that captures the probability that an individual or a system will survive beyond a specific time. One can make assumptions about the shape of the survival distribution. Parametric maximum likelihood methods can then be used to fit the data to the chosen distribution. Rudnick et al. (2016) used exponential distribution to model Spray glider reliability but other models exist, such as LogNormal and that devised by Weibull.

Nonparametric methods can also be used to estimate the probability of survival for a population, without making any assumption with respect to the shape of the distribution (Kalbfleisch and Prentice 2002). A number of studies have used nonparametric methods for estimating the survivability of autonomous underwater

vehicles (Griffiths et al. 2003; Podder et al. 2004; Brito et al. 2014; Rudnick et al. 2016), including the Kaplan–Meier estimator (Kaplan and Meier 1958).

These studies depend on the use of mission data collected in the target operational environment. For missions in extreme environments, such as under ice, there is a paucity of mission data collected from the target environment, making it impossible to use conventional survival methods to estimate the mission risk. To address this problem, Brito et al. (2010) developed an extended version of the Kaplan–Meier estimator. Their estimator uses mission data collected in a benign environment and expert subjective judgment on the impact of those data in the targeted extreme environment.

The extended Kaplan–Meier survival estimator, $\hat{S}$, for quantifying the probability of survival with distance x is presented in Eq. (1) below.

A mission (either failed or successful) is considered as an event. All events are assigned the decreasing index n_i according to the mission distance at which it ended (regardless of the outcome). For each fault, F_i, a group of experts is asked to agree on the probability of fault leading to AUV loss, given that it is operated in a target environment E. This is the probability of failure; it is a conditional probability and it is written as $P(L|F_i,E)$,

$$\hat{S}(x) = \prod_{x_i<x}\left[1 - \left(\frac{1}{n_i}\right)P(L|F_i, E)\right]. \tag{1}$$

This estimator was used to inform operational decision-making for AUV deployments in extreme environments (Brito et al. 2010, 2012). There are two types of risk mitigation that lead to the reduction of the likelihood of failure: 1) monitoring distance and 2) failure prevention and reduction. Details of each are presented in the following subsections.

a. Mitigation with monitoring distance

An important feature of using a survival profile is that it allows us to quantify the impact of implementing a monitoring distance. The engineering purpose is to identify and fix any failures that emerge at short distances. Mathematically, consider that the aim is to travel distance r and that a monitoring distance d is put in place; the conditional survival distribution provides a probability of loss for the target distance r, $P(x < r)$, given that the vehicle has traveled the monitoring distance d. The probability of losing the AUV for a mission with r, given the implementation of d, is then

$$P(x<r|x>d) = \frac{P(x<r) - P(x<d)}{1 - P(x<d)}, \tag{2}$$

where $P(x < r \mid x > d)$ represents the conditional probability of loss in mission up to r given that it has survived monitoring distance up to d, where $d < r$. The expression $P(x < r)$ represents the probability of vehicle loss up to r, and $P(x < d)$ represents the probability of vehicle loss up to d. These probabilities are computed for the survival distribution function. The implementation of a monitoring distance forms a key risk management strategy for AUV missions in critical environments (Griffiths et al. 2003).

The decision to identify the most suitable monitoring distance is informed by both the slope of the survival distribution and the practicality and cost of its implementation. The cost and the practical challenges of implementing the monitoring distance are not discussed in this paper.

With respect to the survival distribution, if the slope of the survival distribution is constant with the distance, then there is no gain in survival by implementing the monitoring distance. On the other hand, if the survival profile shows a steep slope in the first tens of kilometers and then it plateaus for greater distances, then there is a benefit to be gained from implementing the monitoring distance. The optimum monitoring distance is that distance in the survival profile where the survival profile becomes closest to flat.

Depending on the environment, the monitoring distance can be easy or hard to be implemented. The monitoring mission must allow operators to test the functionality of the AUV. Therefore, the AUV must be within communication range of the pilot or control team. The range is also important in terms of recovery. The implementation of a monitoring mission implies that it is possible to recover the AUV at any time if a fault has occurred that needs mending prior to committing to the main missions. Different environments affect the ability to communicate with, or to recover, the AUV. The implementation of the monitoring distance must be defined with an understanding of these constraints.

b. Failure prevention and correction

Failure correction, the process of annulling a failure, involves understanding the failure and putting into place an action to fix it. For AUV risk analysis, failure mitigation was considered in the analysis presented in Brito et al. (2010, 2012) for propelled AUVs and in Brito et al. (2014) and Rudnick et al. (2016) for underwater gliders. There is always a degree of subjective uncertainty as to whether a failure has been mitigated. Brito et al. (2012) capture this uncertainty in the form of a probability of failure mitigation. The authors use the probability of

failure mitigation, elicited from a panel of experts, to update the survival profile. The probability of loss for a given failure, in a given environment, given a mitigation strategy M_i is calculated using

$$P(L|F_i, E, M_i) = P(L|F_i, E)(1 - P_{M_i}), \quad (3)$$

where P_{M_i} is the probability of failure being mitigated. A P_{M_i} value of 1 means that the mitigation action completely mitigates the failure and 0 means that the mitigation does not mitigate the failure. The risk profile calculated using the survival estimator presented in Eq. (1) does not take into account the probability of mitigation. To account for the probability of mitigation, $P(L|F_i, E)$ in Eq. (1) must be replaced by $P(L|F_i, E, M_i)$, which is calculated in Eq. (3).

3. ISE Explorer case study

The ISE Explorer is an autonomous underwater vehicle developed by ISE, Port Coquitlam, British Columbia, Canada. The AUV has a length of 7.4 m, a body diameter of 0.74 m, and is depth rated to 5000 m. The weight of the AUV varies from one mission to another, depending on the payload and battery configuration; for the 2010 and 2011 Arctic campaigns, the weight was 1870 kg. The propulsion is by a propeller with energy for propulsion, controls, and communication from lithium-ion Exide batteries, 30 modules each of 1.6-kWh energy. The maximum range is 450 km at 1.5 m s^{-1} (Kaminski et al. 2010; Crees et al. 2010).

In this case study, we consider the operational data gathered for vehicle B05 during the Arctic operations in 2010 and 2011. The initial risk analysis for the operations in the Arctic was presented in Brito et al. (2012).

The dataset consisted of 32 missions; the fault data are presented in Table 1. For each fault, experts were asked to assess the likelihood of a fault leading to vehicle loss and the likelihood of a failure being mitigated in light of the mitigation action discussed with the engineering team. The expert judgments were elicited at two separate workshops. The first workshop was held in Halifax, Nova Scotia, Canada, from 8 to 10 December 2009; the second workshop was held in Vancouver, British Columbia, Canada, in 2011. The deployments within the dataset were from the following events:

- Fabrication and assembly (May–September 2009)
- Builder sea trials, (8 September–12 October 2009)
- First homing and positioning trials (16 November–4 December 2009)
- Second homing and positioning trials (4 January–28 January 2010)
- Mission testing (22 February–12 March 2010)
- Arctic survey (4 May–22 May 2010)
- Vancouver trials (17 February–22 February 2011)
- Bedford Basin trials (14 June–15 June 2011)

The full fault description and mission details are provided in Brito et al. (2012). In the same paper, the authors present the results of the expert judgment elicitation. A formal expert judgment elicitation was conducted in order to elicit from a group of five experts two risk assessments for each fault. First, for each fault, the experts were asked to agree on the probability of the fault leading to vehicle loss. Following the completion of this process, the experts were then asked to agree on the probability that the failure mitigation strategy proposed by the engineers would correct the failure. Brito et al. (2012, p. 1693) present a table of the agreed expert assessments for all 51 failures in the dataset. When the authors plotted the density distribution of the probability of failure mitigation, they realized that there were three distinct modes in this distribution. The first mode, at zero, comprised assessments for which the failure had not been understood and for which a mitigation plan had not been developed. The second distribution, with mode at 0.5, comprised failures where although a mitigation strategy had been developed, it had not been tested. A third distribution, with mode at 0.9, comprised failure for which a mitigation plan had been developed and tested.

4. Method for updating risk based on mitigation effectiveness

Previous research has assumed that the probability of failure mitigation was a fixed value (Brito et al. 2010, 2012). Once agreed in the workshop, by a group of experts, the assumption is that is value remains constant. Our argument now is that in reality each mission is a test for the mitigation action or strategy. Therefore, the result of the test can be used to update the probability that the failure was mitigated. This can be modeled using a Bayesian theory that captures the rationale that belief in a hypothesis is influenced by new observations (or evidence). The posterior parameter distribution $p(\theta|D)$ is inferred from the observed data D. The prior distribution $p(\theta)$ represents any known information regarding θ, before D is observed.

Before any AUV missions take place, experts agree on the probability of failure mitigation. The probability of failure mitigation tends to be specified as a single probability value, from 0 to 1, rather than a probability density function. The probability of mitigation agreed at the workshop was the prior of the probability of failure

TABLE 1. Probability of failure mitigation for all faults presented in Brito et al. (2012). Column 1 is the fault reference number. Column 2 contains the values of P_{M_i} for each failure. Column 3 presents the probability of loss given the fault, without considering the mitigation. Column 4 presents the number of times that the fault has reoccurred. Columns 5 and 6 present the hyperparameters of the beta distribution fitted to prior P_{M_i}. Columns 7–10 present the properties of the posterior P_M. These columns present the hyperparameters, the mean, and the standard deviation (std dev). Column 11 presents the probability of loss given the fault and the mitigation, taking into account the prior P_{M_i}. Column 12 presents the probability of loss given the fault and the mitigation, considering the posterior P_{M_i}.

Fault ref. No.	P_{M_i}	P(loss|fault)–95% quantile	No. of times fault reoccurred.	Prior		Posterior				With mitigation 95% quantile	Bayesian updated 95% quantile
				a	b	a	b	Mean	Variance		
1	0	0.000 413	1	—	—	—	—	—	—	4.13×10^{-4}	4.13E-04
2	0.9	6.4×10^{-7}	0	12.451	1.383	24.452	1.384	0.946	0.001 89	6.40×10^{-8}	3.46E-08
3	0.8	0.0536	0	20.299	5.075	32.298	5.075	0.864	0.003 06	1.07×10^{-2}	7.29E-03
4	0.9	6.4×10^{-7}	0	12.451	1.383	24.452	1.384	0.946	0.001 89	6.40×10^{-8}	3.46E-08
5	0.95	6.4×10^{-7}	0	6.489	0.341	18.489	0.341	0.982	0.000 898	3.20×10^{-8}	1.15E-08
6	1	6.4×10^{-7}	0	—	—	—	—	—	—	0	0
7	1	0	0	—	—	—	—	—	—	0	0
8	0.9	0.48	0	12.451	1.383	24.452	1.384	0.946	0.001 89	4.80×10^{-2}	2.59E-02
9	0.1	0.004 64	0	3.575	32.175	15.575	32.175	0.326	0.004 51	4.18×10^{-3}	3.13E-03
10	0.95	1	0	6.489	0.341	18.489	0.341	0.982	0.000 898	5.00×10^{-2}	1.80E-02
11	0.95	6.4×10^{-7}	0	6.489	0.341	18.489	0.341	0.982	0.000 898	3.20×10^{-8}	1.15E-08
12	1	0.504	0	—	—	—	—	—	—	0	0
13	0.8	0.0361	0	20.299	5.075	32.298	5.075	0.864	0.003 06	7.22×10^{-3}	4.91E-03
14	0.75	0.921	0	22.43	7.477	34.43	7.477	0.823	0.003 42	2.30×10^{-1}	1.63E-01
15	0.4	0.901	0	45.493	68.239	57.492	68.239	0.457	0.001 96	5.41×10^{-1}	4.89E-01
16	0.95	0.396	0	6.489	0.341	18.489	0.341	0.982	0.000 898	1.98×10^{-2}	7.13E-03
17	1	0	0	—	—	—	—	—	—	0	0
18	0.9	0.0166	0	12.451	1.383	24.452	1.384	0.946	0.001 89	1.66×10^{-3}	8.96E-04
19	0.8	0	0	20.299	5.075	32.298	5.075	0.864	0.003 06	0	0
20	0.4	0.79	0	45.493	68.239	57.492	68.239	0.457	0.001 96	4.74×10^{-1}	4.29E-01
21	0.8	0.0361	0	20.299	5.075	32.298	5.075	0.864	0.003 06	7.22×10^{-3}	4.91E-03
22	0	1	0	—	—	—	—	—	—	1	1
23	0.9	0	0	12.451	1.383	24.452	1.384	0.946	0.001 89	0	0
26	0.95	0.0361	0	6.489	0.341	18.489	0.341	0.982	0.000 898	1.81×10^{-3}	6.50E-04
28	0.95	0.167	0	6.489	0.341	18.489	0.341	0.982	0.000 898	8.35×10^{-3}	3.01E-03
29	0.9	0	0	12.451	1.383	24.452	1.384	0.946	0.001 89	0	0
30	0.8	0.176	0	20.299	5.075	32.298	5.075	0.864	0.003 06	3.52×10^{-2}	2.39E-02
31	0.95	0.0361	0	6.489	0.341	18.489	0.341	0.982	0.000 898	1.81×10^{-3}	6.50E-04
32	1	0	0	—	—	—	—	—	—	0	0
33	1	0	0	—	—	—	—	—	—	0	0
34	0.9	0.009 83	0	12.451	1.383	24.452	1.384	0.946	0.001 89	9.83×10^{-4}	5.31E-04
35a	1	0	0	—	—	—	—	—	—	0	0
35b	0.6	1	0	68.239	45.493	80.239	45.493	0.638	0.001 82	4.00×10^{-1}	3.62E-01
36	0.9	0.798	0	12.451	1.383	24.452	1.384	0.946	0.001 89	7.98×10^{-2}	4.31E-02
37	0.1	0.004 64	0	3.575	32.175	15.575	32.175	0.326	0.004 51	4.18×10^{-3}	3.13E-03
38	1	0	0	—	—	—	—	—	—	0	0
39	0.75	0.109	0	22.43	7.477	34.43	7.477	0.823	0.003 42	2.73×10^{-2}	1.93E-02
40	0.5	0.202	0	59.256	59.256	71.256	59.256	0.546	0.001 89	1.01×10^{-1}	9.17E-02
41	0.5	0.0189	0	59.256	59.256	71.256	59.256	0.546	0.001 89	9.45×10^{-3}	8.58E-03
42	0.5	0.0197	2	59.256	59.256	69.256	61.256	0.531	0.001 89	9.85×10^{-3}	9.24E-03
43	0.5	1	0	59.256	59.256	71.256	59.256	0.546	0.001 89	5.00×10^{-1}	4.54E-01
44	0.5	0.0197	2	59.256	59.256	69.256	61.256	0.531	0.001 89	9.85×10^{-3}	9.24E-03
45	0	0.0191	0	—	—	—	—	—	—	1.91E-02	1.91E-02
46	0.5	0.000 09	2	59.256	59.256	69.256	61.256	0.531	0.001 89	4.50×10^{-5}	4.22E-05
47	0.5	0.78	0	59.256	59.256	71.256	59.256	0.546	0.001 89	3.90×10^{-1}	3.54E-01
48	0.5	0.451	0	59.256	59.256	71.256	59.256	0.546	0.001 89	2.26×10^{-1}	2.05E-01
50	0.5	1	0	59.256	59.256	71.256	59.256	0.546	0.001 89	5.00×10^{-1}	4.54E-01
51	0.95	0.0361	0	6.489	0.341	18.489	0.341	0.982	0.000 898	1.81×10^{-3}	6.50E-04
52	0.5	0.78	0	59.256	59.256	71.256	59.256	0.546	0.001 89	3.90×10^{-1}	3.54E-01
53	0.5	0.78	0	59.256	59.256	71.256	59.256	0.546	0.001 89	3.90×10^{-1}	3.54E-01
54	0.1	0.009 47	1	3.575	32.175	14.575	33.175	0.305	0.004 35	8.52×10^{-3}	6.58E-03

mitigation that must be updated in light of successful missions as well as reoccurrences of failures.

There are two key problems. First, we must model the prior probability of failure mitigation in such a way that allows us to update its value based on field observations. Second, we must have means to conduct the Bayesian inference. This is discussed in sections 4a and 4b.

a. Modeling the prior

In the process of building the risk model, the experts agree on the probability that the failure correction action will remove the failure. In Brito et al. (2012), this is denoted as the probability of failure mitigation and it is represented by P_{M_i}. In this paper our aim is to update P_{M_i} in light of subsequent missions. The quantity P_{M_i} is the a priori probability of mitigation, which we aim to update using Bayesian inference.

There is uncertainty associated with the estimate of P_{M_i}. This uncertainty was not captured in the expert judgment elicitation presented in Brito et al. (2012). Nevertheless, similar to the way that there is uncertainty associated with the probability of loss, there is also uncertainty associated with the probability of mitigation. In addition, P_{M_i} must be modeled in a way that allows us to apply Bayesian inference. To enable these two steps, the beta probability density function (pdf) is selected for two reasons. First, it is the most suitable probability distribution to model expert judgments for single-mode assessments (O'Hagan et al. 2006). Second, this distribution is a conjugate distribution for the binomial distribution. This is to say, if the beta distribution is used as a prior pdf of the probability of failure mitigation and the conditional distribution is binomial, then the posterior is always a beta distribution.

The probability of failure mitigation θ_i is taken to follow the beta distribution, which is

$$p(\theta) = \frac{1}{B(a,b)}\theta_i^{a-1}(1-\theta_i)^{b-1}, \tag{4}$$

where a and b are the constant hyperparameters of the beta distribution, and $B(a, b)$ represents the beta function,

$$B(a,b) = \int_0^1 \theta^{a-1}(1-\theta)^{b-1}\,d\theta. \tag{5}$$

The beta function is the normalization constant for the beta distribution.

The hyperparameters of the beta distribution are calculated using P_{M_i}. Equations (6) and (7) are obtained by manipulating the equations for the mean and variance of the beta distribution,

$$a_j = \frac{\mu_j^2}{\sigma_j^2} - \frac{\mu_j^3}{\sigma_j^2} - \mu_j, \quad \text{and} \tag{6}$$

$$b_j = \frac{a_j}{\mu_j} - a_j. \tag{7}$$

Both the mean μ_j and the variance σ_j^2 for each failure j are obtained from the prior assessments of the probability of failure mitigation. The mean equals the P_{M_i} estimated by the experts. The variance for each probability of failure mitigation is calculated from the trimodal probability of mitigation distribution (Brito et al. 2012). The details of the characteristics of these modes are presented in section 5a. The variance for each mode was calculated using

$$\sigma_y^2 = \sum_{k=1}^{m}(x_k - \mu_y)^2 p_k, \tag{8}$$

where y is the mode number, 1–3. The index $k = 1, m$ is the number of probability classes in each mode of the trimodal probability of mitigation distribution, and p_k is the proportion of assessments in each probability class. For example, for mode 1, $p_1 = 0.571$ and $p_2 = 0.429$. Variable x_k is the probability associated with each class, and μ_y is the mean probability of each mode. In this case the variance for mode 1 is 0.002 449, for mode 2 it is 0.002 09, and for mode 3 it is 0.006 06. The variance for each failure σ_j^2 is equal to the variance for the mode y that encapsulates this failure.

b. Bayesian inference

Having defined the prior for the probability of failure mitigation, the next step is to update this prior in light of subsequent missions. We consider that each mission that the AUV conducts is a test of the effectiveness of the mitigation action. The probability of failure mitigation is analogous to the probability of success in a binomial trial. Each mission is a trial, which is successful if the mission was completed failure free and unsuccessful if the failure in question occurred during the mission. The total number of successful missions is denoted as m. If we denote the total number of missions as n, with θ being the probability of success, and $p(\theta)$ the probability distribution of θ, then the probability of success for m out of n missions is calculated using the binomial expression

$$P(m|\theta,n) = \binom{n}{m} \times \theta^m \times (1-\theta)^{n-m}. \tag{9}$$

The same convention is used as in section 4a because θ is the probability of failure mitigation.

The aim is to estimate the value of θ given the observations made with respect to n and m. To achieve this aim,

we must calculate $P(\theta|m,n)$. The Bayesian rule allows us to calculate this probability using Eq. (10). This equation is demonstrated from first principles in the appendix,

$$P(\theta|m,n) = \frac{P(m|\theta,n) \times P(\theta)}{P(m|n)}, \tag{10}$$

where $P(m|n)$ is the normalization constant, uniquely defined by requiring the total posterior probability to be 1. It is the probability of a successful trial given that a number of tests are conducted. Regardless of whether we ignore $P(m|n)$, we can argue that $P(\theta|m, n)$ is proportional to the numerator of Eq. (11). Equation (11) is the beta-binomial inference for the probability θ given m experiments and n successes. The expressions for $P(m|\theta, n)$ and $p(\theta)$ are given in Eqs. (9) and (4), respectively,

$$\begin{aligned} p(\theta|m,n) &\propto P(m|\theta,n) \times p(\theta) \\ &\propto \binom{n}{m} \theta^r \times (1-\theta)^{n-m} \times \frac{1}{B(a,b)} \times \theta^{a-1} \times (1-\theta)^{b-1} \\ &\propto \binom{n}{m} \times \frac{1}{B(a,b)} \times \theta^{m+a-1} \times (1-\theta)^{n-m+b-1} \\ &\propto \text{Beta}(m+a, n-m+b), \end{aligned} \tag{11}$$

where $P(\theta|m, n)$ is proportional to a beta distribution. We can use Eq. (11) to calculate the updated probability of failure mitigation.

5. Bayesian updated judgments

To illustrate the application of the Bayesian inference approach, we applied the approach to the ISE Explorer AUV B5. Twelve missions were performed after the failure mitigation assessment process in January 2010. Of these, six missions—dives 51–56— were carried out in the Arctic. Four missions took place off from Vancouver on 17, 18, 21, and 22 February 2011, and two missions were carried out in the Bedford Basin on 14 and 15 June 2011. Here we applied Bayesian inference for updating the likelihood that failure x had been mitigated.

a. Updating failure risk

Figure 1 presents the cumulative distribution functions (CDFs) for the probability of failure mitigation for failures 9, 13, 15, 40, 42, and 35b. These are typical examples of the probability of mitigation for the failures in the three modes of the probability of failure mitigation, for where failures did and did not occur.

The posterior probability of failure mitigation increases even if one or two failures emerge during subsequent trials. This suggests that the failures are rarer than the prior (expert) probabilities suggest.

The updated values for all probabilities of mitigation are presented in Table 1. Where the prior P_{M_i} is 0 or 1 it is impossible to fit a beta distribution and thus it is not possible to update these probability of mitigation estimates.

The table shows that most failures did not reemerge in subsequent missions.

Figure 2 presents a summary of the effect of the probability of mitigation, before and after subsequent missions, for failures for which the probability of loss was greater than 0.01.

In Brito et al. (2012) the authors show that the distribution of the probability of mitigation can be trimodal. Figure 3 shows the a priori pdf of the probability of mitigation, with no prior knowledge of the effectiveness of the failure mitigation (gray) and with knowledge of subsequent missions (black). The three modes of the distribution identified in Brito et al. (2012) (gray) are still evident in the a posteriori distribution (black).

b. Reliability growth

Having calculated the posterior for the probability of failure mitigation (see Table 1), it is then possible to calculate the updated survival profile for the autonomous underwater vehicle using the extended version of the Kaplan–Meier estimator, Eq. (1). Figure 4 shows the survival distribution.

Figure 4 shows that for long missions, the largest increase in survivability comes from addressing the historical failures with their a priori estimated probability of effective mitigation. Considering the survival profiles for mitigated and Bayesian updated, for a mission between 57 and 324 km, the probability of survival increased by 1.6% (see Fig. 4). These results are based on the mean estimate for the updated probability of mitigation for each failure. From Table 1, it is possible to see that there is a reduction in the variance from the prior estimate for the probability of mitigation and the posterior estimate. The results therefore show that with the Bayesian inference, there is an increase in the probability of survival and an increase in confidence.

From the survival distribution for the unmitigated case, it is possible to see that the probability of survival decreases by approximately 15% in the first 31 km of a mission. This is the most significant slope in the survival distribution before it plateaus and it informed the implementation of mission 51. Two other test missions

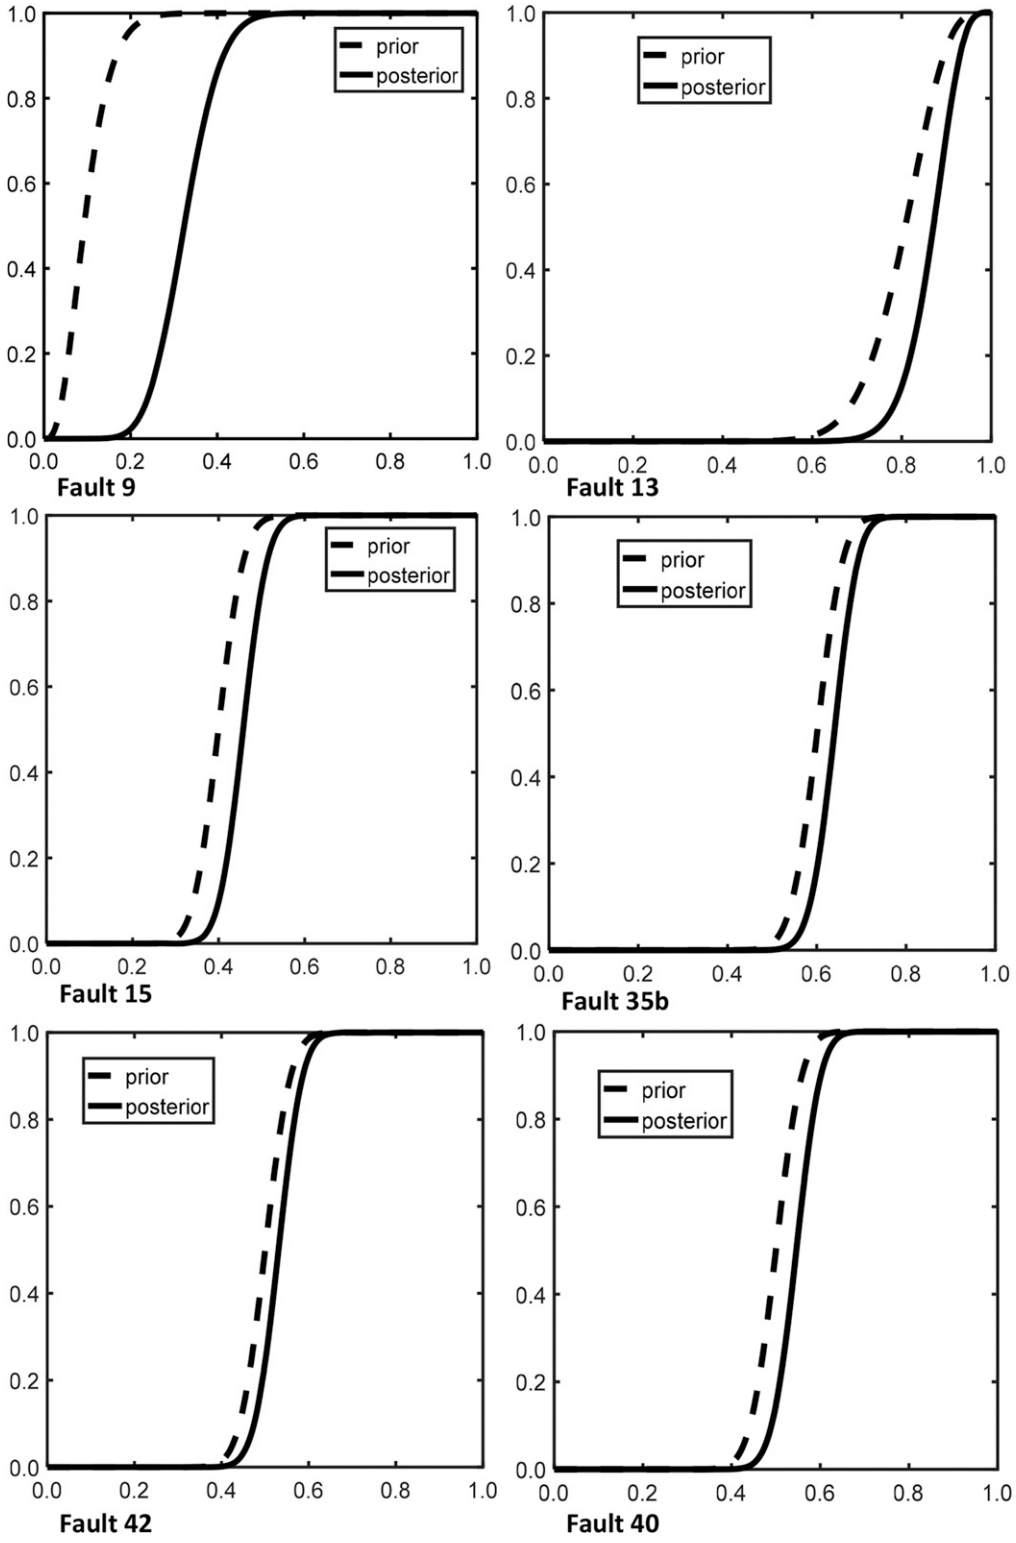

FIG. 1. Density distribution for the mean of the probability of fault mitigation for failures 9, 13, 15, 35b, 42, and 40.

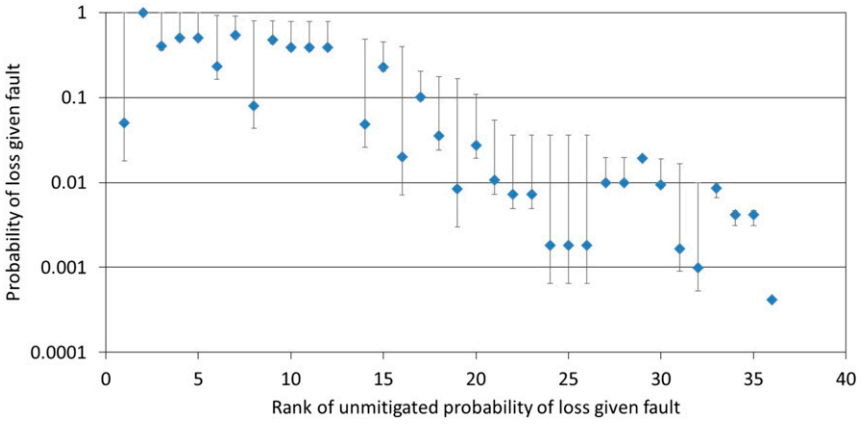

FIG. 2. Probability of loss given fault, mitigation, and trials results. Ranked order by unmitigated probability of loss given fault, for those faults above 0.01. The upper point of the "error bar" is the unmitigated data, the main point is after a priori estimated mitigation, and the lower error bar is posterior after Bayesian inference.

were conducted, missions 52 and 53. For missions 51 and 52, the vehicle was constantly monitored using short-range localization (SRL). Table 2 presents the missions conducted by the ISE Explorer vehicle B05 in the Arctic during the campaign in 2010.

Table 3 presents the survival estimates for missions 51–53, considering the unmitigated, mitigated, and Bayesian updated survival profiles. The probability of survival for missions 52 and 53, given the implementation of a monitoring distance of 31 km, was calculated using Eq. (2).

The survivability of the AUV, for mission 51, increased by 0.06, from the mitigated survival profile to the Bayesian updated survival profile. For missions greater than 31 km, the probability of survival for the Bayesian updated risk profile is 0.116 higher than the probability of survival estimated using the mitigated survival profile.

Table 4 presents the survival estimates for missions 55 and 56. The survival profile contains data up to a distance of 334 km; therefore, the probability of survival for mission 54 cannot be calculated. The product rule was used to estimate the overall probability of surviving both survey missions.

Considering no failure mitigation, for the case where a monitoring distance of 87 km is implemented (missions 51 and 52), the probability of surviving missions 55 and 56 increases by 17.7%. On the other hand, if we consider

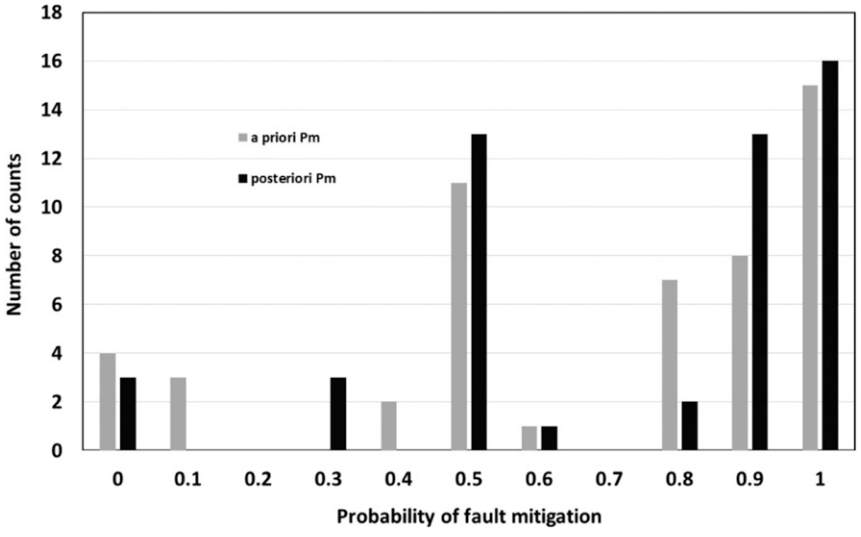

FIG. 3. Probability of failure mitigation: the prior probability of mitigation distribution (gray) and the posterior probability of mitigation distribution (black).

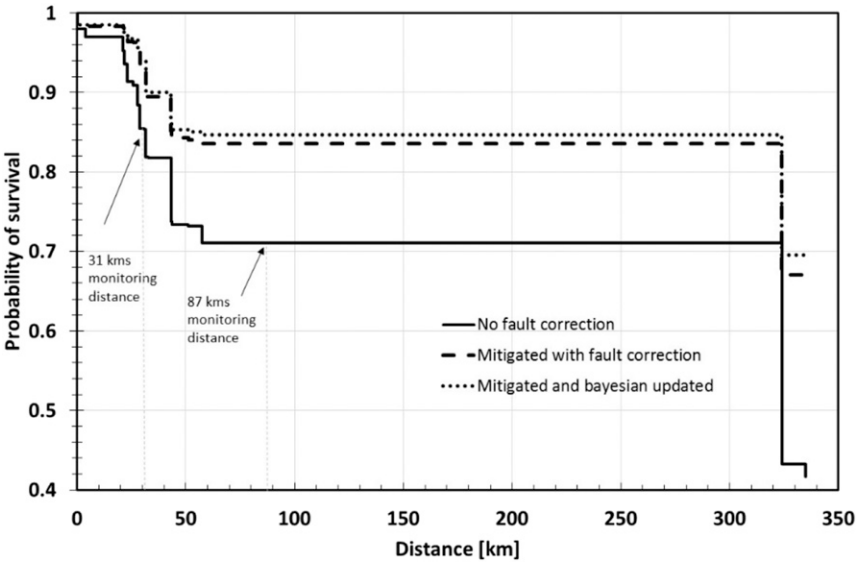

FIG. 4. Survival distribution for the ISE Explorer vehicle deployment in the Arctic. The survivability without mitigation (sold line), the survivability considering the mitigation based on the prior belief (dashed line), and the survivability considering the mitigation updated (dotted line).

only failure correction and the Bayesian inference, with no monitoring distance, the probability of survival increases by 26.2%. This shows the importance of quantifying the impact of failure correction and of updating the risk profile based on subsequent missions.

6. Conclusions

Autonomous underwater vehicles are complex systems where managing the risk of premature abort or loss is very important. While designers and manufacturers will have sought to design and build reliable components, hardware subsystems, and software, a vehicle in the hands of users who were not the developers is still likely to show faults and failures. Developers and experts are efficient at managing risk through mitigation—that is to say, describing and understanding the failures that emerge, coming up with solutions, and then testing those solutions. However, this is usually a subjective process without an analysis of the likely impact of the fault or the effectiveness of failure mitigation on the probability of a vehicle surviving a mission.

This paper injects transparency into this process. It argues that while an initial assessment can be obtained for the probability of successful mitigation for a failure, the actual probability of mitigation must be updated based on the results of subsequent missions.

Tracking reliability growth is required in order to ensure effective outcomes from the deployment of autonomous systems. In this paper we present Bayesian formalism for tracking the reliability growth of autonomous underwater vehicles. Each mission was considered as a test in a binomial trial. We applied the method to update the risk of the ISE Explorer AUV, following the Arctic campaign in 2010.

Another potential application of this approach within the AUV context is the issue of updating the availability prediction. In an AUV deployment there are several phases, for example, pretest, posttest, vehicle overboard (if being launched from a ship), and so on. Availability is the likelihood of a successful sequence of phases taking place (Brito and Griffiths 2011). In actual field deployment, the transition from one phase to the next can be considered to be a binomial trial. Similar to what has been presented in this paper, the a priori of the probability of successful transition can be updated using the results of field deployments, thus allowing us to update the availability prediction of an AUV.

In our view this study has highlighted some limitations in the expert judgment elicitation, as the practice of eliciting a single value for the probability of mitigation

TABLE 2. All missions conducted by ISE Explorer B05 in the Arctic campaign in 2010.

Mission	Distance (km)
51	31
52	56
53	131
54	336
55	326
56	325

TABLE 3. Survival estimates for ISE Explorer's missions 51–53. The monitoring mission of 31 km corresponds to mission 51.

Scenario	Probability of survival			Survival for all missions
	Mission 51	Mission 52	Mission 53	
Unmitigated	0.854	0.732	0.710	0.444
Mitigated	0.934	0.84	0.835	0.655
Bayesian updated	0.94	0.851	0.851	0.681
Mitigated + monitor 31 km		0.899	0.894	0.8048
Bayesian updated + monitor 31 km		0.905	0.905	0.8201

presents a problem. Experts tend to assign a probability of 0 if the event is very unlikely and a probability of 1 if it is very likely. In the elicitation of the probability of loss given a failure, the elicitation process encourages the experts to assign a probability distribution, such as a beta distribution. A similar approach should be adopted for eliciting the probability of failure mitigation. By doing so it allows the decision-maker to update the risk for those failures in light of subsequent missions. This research shows that the practice of eliciting a single probability value for the probability of failure mitigation can be problematic. It forces the analyst to make assumptions regarding modeling the probability of mitigation. In making these assumptions, the analyst may introduce bias. This is evident with Fig. 3; here the probability of failure mitigation is represented as the number of counts. The probability of failure mitigation provided by experts was discrete values 0, 0.1, ... , in increments of 0.1. This allowed us to model the probability of failure mitigation as a number of counts. If a more detailed probability of failure mitigation had been obtained, instead of using the number of counts to represent the probability of failure, then mitigation could have used as density function. In the case study presented in this paper, we modeled the probability of failure mitigation with a beta distribution. To match a beta distribution to the probability of failure mitigation, we assumed that the mean of the beta distribution was the probability of failure mitigation. The width of the prior (expert) probability distributions is also not known. To model the width, we assumed that the prior variance was the same as the variance for the mode of the priori probability of failure mitigation distribution; this was calculated using Eq. (7). Alternatively, we could have assumed that the mode of the beta distribution was the probability of failure mitigation or have made another assumption for the variance. This is, in our view, one limitation of this research. This highlights the need for analysts to elicit a probability distribution for the probability of failure mitigation. To our knowledge this is not current practice in risk modeling.

Acknowledgments. The authors thank the anonymous reviewers for their very insightful comments. We also thank the experts who took part in the expert judgment elicitation conducted in Halifax and in Vancouver. This work was partly funded by the Natural and Environment Research Council Grant NE/I015647/1 awarded to Mr. Gwyn Griffiths as Principal Investigator.

APPENDIX

Derivation of the Bayesian Equation for Three Variables

Let $\boldsymbol{\theta}$ be a vector of probabilities of success for several independent binomial trials. Let $\mathbf{n}$ be a vector of the number of trials and let $\mathbf{m}$ be a vector of the number of successes.

The joint distribution $P(\boldsymbol{\theta}, \mathbf{m}, \mathbf{n})$ can be calculated using

$$\begin{aligned} P(\boldsymbol{\theta}, \mathbf{m}, \mathbf{n}) &= P(\boldsymbol{\theta}|\mathbf{m}, \mathbf{n}) \times P(\mathbf{m}, \mathbf{n}) \\ &= P(\boldsymbol{\theta}|\mathbf{m}, \mathbf{n}) \times P(\mathbf{m}|\mathbf{n}) \times P(\mathbf{n}), \quad \text{and} \\ P(\boldsymbol{\theta}, \mathbf{m}, \mathbf{n}) &= P(\mathbf{m}|\boldsymbol{\theta}, \mathbf{n}) \times P(\boldsymbol{\theta}, \mathbf{n}) \\ &= P(\mathbf{m}|\boldsymbol{\theta}, \mathbf{n}) \times P(\boldsymbol{\theta}|\mathbf{n}) \times P(\mathbf{n}) \\ &= P(\mathbf{m}|\boldsymbol{\theta}, \mathbf{n}) \times P(\boldsymbol{\theta}) \times P(\mathbf{n}), \end{aligned}$$

TABLE 4. Survival estimates for ISE Explorer's Arctic survey missions.

	Survival estimates for survey missions		Survival for all missions
	Mission 55	Mission 56	
Unmitigated	0.433	0.433	0.187
Mitigated	0.67	0.67	0.449
Bayesian updated	0.695	0.695	0.483
Unmitigated + monitor 87 km	0.610	0.610	0.372
Mitigated + monitor 87 km	0.802	0.802	0.644
Bayesian updated + monitor 87 km	0.817	0.817	0.667

where $\boldsymbol{\theta}$ is independent from $\mathbf{n}$ because we must know both $\mathbf{m}$ and $\mathbf{n}$ in order to infer $\boldsymbol{\theta}$. Thus, $\mathbf{P}(\boldsymbol{\theta}|\mathbf{n}) = \mathbf{P}(\boldsymbol{\theta})$.

Taking into account the two terms for $P(\boldsymbol{\theta}, \mathbf{m}, \mathbf{n})$,

$$P(\boldsymbol{\theta}|\mathbf{m},\mathbf{n}) \times P(\mathbf{m}|\mathbf{n}) \times P(\mathbf{n}) = P(\mathbf{m}|\boldsymbol{\theta},\mathbf{n}) \times P(\boldsymbol{\theta}) \times P(\mathbf{n}),$$

and

$$P(\boldsymbol{\theta}|\mathbf{m},\mathbf{n}) = \frac{P(\mathbf{m}|\boldsymbol{\theta},\mathbf{n}) \times P(\boldsymbol{\theta})}{P(\mathbf{m}|\mathbf{n})}.$$

REFERENCES

Brito, M. P., and G. Griffiths, 2011: A Markov chain state transition approach to establishing critical phases for AUV reliability. *IEEE J. Oceanic Eng.*, **36**, 139–149, https://doi.org/10.1109/JOE.2010.2083070.

——, ——, and P. Challenor, 2010: Risk analysis for autonomous underwater vehicle operations in extreme environments. *Risk Anal.*, **30**, 1771–1788, https://doi.org/10.1111/j.1539-6924.2010.01476.x.

——, ——, J. Ferguson, D. Hopkin, R. Mills, R. Pederson, and E. MacNeil, 2012: A behavioral probabilistic risk assessment framework for managing autonomous underwater vehicle deployments. *J. Atmos. Oceanic Technol.*, **29**, 1689–1703, https://doi.org/10.1175/JTECH-D-12-00005.1.

——, D. A. Smeed, and G. Griffiths, 2014: Underwater glider reliability and implications for survey design. *J. Atmos. Oceanic Technol.*, **31**, 2858–2870, https://doi.org/10.1175/JTECH-D-13-00138.1.

BSI, 1991: Quality vocabulary—Part 3: Availability, reliability and maintainability terms; Section 3.1—Guide to concepts and related definitions. British Standards Institution Publ. BS 4778-3.1:1991, 32 pp.

Crees, T., and Coauthors, 2010: UNCLOS under ice survey—An historic AUV deployment in the Canadian high arctic. *Proc. OCEANS 2010 MTS/IEEE*, Seattle, WA, IEEE, 8 pp., https://doi.org/10.1109/OCEANS.2010.5664438.

Eriksen, C. C., T. J. Osse, R. D. Light, T. Wen, T. W. Lehman, P. L. Sabin, J. W. Ballard, and A. M. Chiodi, 2001: Seaglider: A long-range autonomous underwater vehicle for oceanographic research. *IEEE J. Oceanic Eng.*, **26**, 424–436, https://doi.org/10.1109/48.972073.

Feather, M. S., and S. L. Cornford, 2003: Quantitative risk-based requirements reasoning. *Requir. Eng.*, **8**, 248–265, https://doi.org/10.1007/s00766-002-0160-y.

Griffiths, G., and A. Trembanis, 2007: Eliciting expert judgment for the probability of AUV loss in contrasting operational environments. *15th International Symposium on Unmanned Untethered Submersible Technology 2007*, Autonomous Undersea Systems Institute, 494–510.

——, N. W. Millard, S. D. McPhail, P. Stevenson, and P. G. Challenor, 2003: On the reliability of the Autosub autonomous underwater vehicle. *Underwater Technol.*, **25**, 175–184, https://doi.org/10.3723/175605403783101612.

——, M. Brito, I. Robbins, and M. Moline, 2009: Reliability of two REMUS-100 AUVs based on fault log analysis and elicited expert judgment. *16th Annual International Symposium on Unmanned Untethered Submersible Technology 2009 (UUST 09)*, Autonomous Undersea Systems Institute, 451–463.

Kalbfleisch, J. D., and R. L. Prentice, 2002: *The Statistical Analysis of Failure Time Data.* 2nd ed. Wiley Series in Probability and Statistics, Wiley, 489 pp.

Kaminski, C., T. Crees, J. Ferguson, A. Forrest, J. Williams, D. Hopkin, and G. Heard, 2010: 12 days under ice—An historic AUV deployment in the Canadian High Arctic. *Proc. 2010 IEEE/OES Autonomous Underwater Vehicles*, Monterey, CA, IEEE, 11 pp., https://doi.org/10.1109/AUV.2010.5779651.

Kaplan, E. L., and P. Meier, 1958: Nonparametric estimation from incomplete observations. *J. Amer. Stat. Assoc.*, **53**, 457–481, https://doi.org/10.1080/01621459.1958.10501452.

O'Hagan, A., C. E. Buck, A. Daneshkhah, J. R. Eiser, P. H. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow, 2006: *Uncertain Judgments: Eliciting Experts' Probabilities.* Statistics in Practice, Wiley, 338 pp.

Podder, T. K., M. Sibenac, H. Thomas, W. Kirkwood, and J. G. Bellingham, 2004: Reliability growth of autonomous underwater vehicle—Dorado. *OCEANS '04: MTS/IEEE Techno-Ocean '04 (OTO'04) Conference Proceedings*, Vol. 2, IEEE, 856–862, https://doi.org/10.1109/OCEANS.2004.1405576.

Reason, J., 1990: *Human Error.* Cambridge University Press, 302 pp.

Rudnick, D. L., R. Davis, and J. T. Sherman, 2016: Spray underwater glider operations. *J. Atmos. Oceanic Technol.*, **33**, 1113–1122, https://doi.org/10.1175/JTECH-D-15-0252.1.

Singh, H., A. Can, R. Eustice, S. Lerner, N. McPhee, O. Pizarro, and C. Roman, 2004: Seabed AUV offers new platform for high-resolution imaging. *Eos, Trans. Amer. Geophys. Union*, **85**, 289–296, https://doi.org/10.1029/2004EO310002.

Subramanian, S., L. Elliott, R. V. Visnuvajjala, W. T. Tsai, and R. Mojdehbakhsh, 1996: Fault mitigation in safety-critical software systems. *Proceedings: Ninth IEEE Symposium on Computer-Based Medical Systems*, IEEE Computer Society Press, 12–17.

Webb, D. C., P. J. Simonetti, and C. P. Jones, 2001: SLOCUM: An underwater glider propelled by environmental energy. *IEEE J. Oceanic Eng.*, **26**, 447–452, https://doi.org/10.1109/48.972077.