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

Mario P. Brito[1]
University of Southampton, Southampton, Hampshire, UK

Gwyn Griffiths
Autonomous Analytics, Southampton, Hampshire, UK

[1] *Corresponding author address*: Mario P. Brito, University of Southampton, University Road, Southampton, SO17 1BJ
E-mail: m.p.brito@soton.ac.uk

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 non-developers. We show that this difference in reliability is due to the failure prevention and correction procedures - risk mitigation - put in place by developers. However, no formalisation has been developed for updating the risk profile based on the expected effectiveness of the failure prevention and correction process. In this paper we present a generic Bayesian approach for updating the risk profile, 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 ISE Explorer AUV.

Keywords: Instrumentation/sensors, Risk assessment, statistical techniques.

1. Introduction

Autonomous underwater vehicles (AUVs) are mechatronic robotic systems able to navigate underwater whilst 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 as 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 non-developers, 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 as failure faults that led to loss of control over buoyancy and vehicle loss. In the study by Brito et al. (2014), the authors considered as failures vehicle loss. Both studies argue that understanding and eliminating failure modes are key to increasing the probability of successful mission completion and of 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 (BS 4778, 1991), 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 which can potentially lead to AUV loss. Generally, failure, or 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 or risk, mitigation is achieved by reducing the likelihood of failure occurrence.

Whilst 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.

83 The novelty of the technology makes it impossible to obtain past data on the probability of
84 failure mitigation. Therefore, the assessment of the probability of failure mitigation for this
85 type of technology must rely on expert subjective judgment. The same approach has been
86 adopted in space exploration (Feather and Cornford, 2003).
87 Feather and Cornford (2003) presented a hazard management framework to monitor and
88 update the likelihood of the occurrence of design failure modes occurring. The system, Defect
89 Detection and Prevention (DDP), is a probabilistic model. The key assumption is that each
90 failure may have a number of prevention, analysis, control and test (PACTs) methods. The
91 efficiency of each PACT in mitigating the failure is assessed by a group of experts in the field.
92 The DDP system considers that multiple PACTs may have adverse or positive effects on a
93 failure mode. PACTs are not independent, and they may introduce a failure into the system.
94 The probability model aggregates all these effects to quantify the likelihood of failure
95 mitigation. Brito et al. (2012) also used probability for modelling failure mitigation. However,
96 in contrast to Feather and Cornford (2003), the mitigation actions were considered to be
97 independent.
98 In this paper, we present a Bayesian approach to updating the risk profile of an AUV, based
99 on the pre-implementation perceived effectiveness of the mitigation by subject matter
100 experts and the observed performance during subsequent missions. We present a case study
101 of the International Submarine Engineering (ISE) Explorer AUV to illustrate the application of
102 the method, based on the initial failure and survivability data presented in Brito et al. (2012).
103 This paper is organized as follows. Section 2 presents a summary of the methods used for
104 estimating AUV survival. Section 3 presents the data of the ISE Explorer campaigns in the
105 Arctic in 2010 and 2011. Section 4 presents the method proposed for updating the risk profile
106 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 total mission time or mission length. In statistical survival modelling, this observation where the end, or the last reading, did not result in death, is denoted as censored. For AUVs, successful missions are modelled as right censored, which means that the vehicle has survived at least the time or distance travelled by the vehicle. Some missions, however, end in failure. In statistical modelling, 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.

Non-parametric 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 non-parametric 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 an paucity of

131 mission data collected from the target environment, making it impossible to use conventional
132 survival methods to estimate the mission risk. To address this problem, Brito et al (2010)
133 developed an extended version of the Kaplan Meier estimator. Their estimator uses mission
134 data collected in a benign environment and expert subjective judgment on the impact of that
135 data in the targeted extreme environment.

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

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

143 $$\hat{S}(x) = \prod_{x_i<x}\left(1 - \left(\frac{1}{n_i}\right)P(L|F_i,E)\right) \qquad (1)$$

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

148

149 **2.1 Mitigation with monitoring distance**

150 An important feature of using a survival profile is that it allows us to quantify the impact of
151 implementing a monitoring distance. The engineering purpose is to identify and fix any
152 failures that emerge at short distances. Mathematically, consider that the aim is to travel
153 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 travelled a distance d. The probability of losing the AUV for a mission with distance r, given the implementation of a monitoring distance d, is then:

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

where P(x<r|x>d) represents the conditional probability of loss in mission up to distance r given that it has survived monitoring distance up to distance d, where d < r. P(x<r) represents the probability of vehicle loss up to distance r, and P(x < d) represents the probability of vehicle loss up to distance 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 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 which 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.

2.2 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) and Brito et al. (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 or not a failure has been mitigated. In Brito et al. (2012) the authors 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 Eq. (3).

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

Where P_{Mi} is the probability of failure being mitigated. P_{Mi} value of 1 means that the mitigation action completely mitigates the failure and 0 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. In order to account for the probability of mitigation $P(L|F_i, E)$ in Eq. (1) must be replaced by $P(L|F_i, E, Mi)$ calculated in Eq. (3).

3. ISE Explorer Case Study

The ISE Explorer is an autonomous underwater vehicle developed by ISE Limited, Port Coquitlam, Canada. The AUV has a length of 7.4m, a body diameter of 0.74m and is depth rated to 5000m. 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 1870kg. The propulsion is by a propeller with energy for propulsion, controls and communication from Li-ion Exide batteries, 30 modules each of 1.6 kWh energy. The maximum range is 450 km at 1.5 m/s. (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 is presented in Table 1, below. 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, in 2011. The deployments within the dataset were from:

- Fabrication and assembly (May - September 2009)
- Builders 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. Firstly, 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, p1693), 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 realised 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 modelled using a Bayesian theory, which 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 mitigation which must be updated in light of successful missions as well as re-occurrences 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 4.1 and 4.2.

4.1 Modelling 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_{Mi}. In this paper, our aim is to update P_{Mi} in the light of subsequent missions. P_{Mi} is the a priori probability of mitigation, which we aim to update using Bayesian inference.

There is uncertainty associated with the estimate of P_{Mi}. 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, the P_{Mi} must be modelled 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. Firstly, it is the most suitable probability distribution to model expert judgments for single mode assessments (O'Hagan et al., 2006).

Secondly, 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 this is.

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

where *a* and *b* are the constant hyper-parameters of the beta distribution and B(a,b) represents the beta function,

$$\mathrm{B(a,b)} = \int_0^1 \theta^{\,\mathrm{a-1}} \times (1-\theta)^{\mathrm{b-1}} \mathrm{d}\,\theta \quad (5)$$

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

The hyper parameters, *a* and *b*, of the beta distribution are calculated using the probability of failure mitigation P_{Mi}. Eq. (6) and Eq. (7), below, are obtained by manipulating the equations for the mean and variance of the beta distribution.

$$a_j = \frac{\mu_j^2}{\sigma_j^2} - \frac{\mu_j^3}{\sigma_j^2} - \mu_j \quad (6)$$

$$b_j = \frac{a_j}{\mu_j} - a_j \quad (7)$$

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 (μ_j) equals the P_{Mi} estimated by the experts. The variance for each probability of failure mitigation is calculated from the tri-modal probability of mitigation distribution Brito et al. (2012). The details of the characteristics of these modes are presented in section 5.1. The variance for each mode was calculated using.

$$\sigma_y^2 = \sum_{k=1}^{m} \left(x_k - \mu_y\right)^2 p_k \quad (8)$$

where y is the mode number, 1 to 3. The index k =1 to m, is the number of probability classes in each mode of the tri modal 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 $p_2 = 0.429$. 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 0.002449, for mode 2 is 0.00209 and for mode 3 is 0.00606. The variance for each failure σ_j^2 is equal to the variance for the mode y that encapsulates this failure.

4.2 Bayesian inference

Having defined the prior for the probability of failure mitigation, the next step is to update this prior in the 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, p(θ) the probability distribution of θ, the probability of success for m out of n missions is calculated using the binomial expression.

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

The same convention is used as in 4.1 because θ is the probability of failure mitigation.

The aim is to estimate the value of θ given the observations made with respect to the number of trials, n, and the number of successes, m. To achieve this aim we must calculate

P(θ /*m,n*). The Bayesian rule allows us to calculate this probability using Eq. (10), below. This equation is demonstrated from first principles in Appendix A.

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

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. Whether or not we ignore the normalisation constant, P(m|n) we can argue that P(θ|m,n) is proportional to the numerator of Eq. (11), below. Equation 11 is the beta-binomial inference for the probability θ given m experiments and n successes. The expressions for P(m|θ,n) and p(θ) are given in Eq. (9) and Eq. (4), respectively.

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

P(θ|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, 52, 53, 54, 55 and 56, were carried out in the Arctic. Four missions took place off 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.

5.1 Updating failure risk

Fig. 1, below 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.

<Fig. 1 goes here>

The updated values for all probabilities of mitigation are presented in Table 1, below. Where the prior P_{Mi} 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 re-emerge in subsequent missions.

361

362 < Table 1 goes here>

363 Fig. 2, below, presents a summary of the effect of the probability of mitigation, before
364 and after subsequent missions, for failures for which the probability of loss was greater than
365 0.01.

366

367 <Fig. 2 goes here>

368 In Brito et al. (2012) the authors show that the distribution of the probability of mitigation
369 can be tri-modal. Fig. 3, below, shows in black the *a priori* pdf of the probability of mitigation,
370 with no prior knowledge on the effectiveness of the failure mitigation. Fig. 3, in grey shows
371 the probability of failure mitigation without knowledge of subsequent field missions. In black
372 is the updated probability of mitigation with knowledge of field missions. The three modes of
373 the distribution, identified by Brito et al. (2012) are still evident for the *a posteriori*
374 distribution.

375

376 <Fig. 3 goes here>

377 **5.2 Reliability Growth**

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

382

383 <Fig. 4 goes here>

384

385 Fig. 4 shows that for long missions the largest increase in survivability comes from
386 addressing the historical failures with their a priori estimated probability of effective
387 mitigation. Considering the survival profiles for mitigated and Bayesian updated, for a mission
388 between 57 and 324 km, the probability of survival increased by 1.6%, (see Fig. 4). These
389 results are based on the mean estimate for the updated probability of mitigation for each
390 failure. From Table 1, it is possible to see that there is a reduction in the variance from the
391 prior estimate for the probability of mitigation and the posterior estimate. The results
392 therefore show that with the Bayesian inference there is an increase in the probability of
393 survival and an increase in confidence.
394 From the survival distribution for the un-mitigated case, it is possible to see that the
395 probability of survival decreases approximately by 15% in the first 31km of a mission. This is
396 the most significant slope in the survival distribution before it plateaus and it informed the
397 implementation of mission 51. Two other test missions were conducted, missions 52 and 53.
398 For missions 51 and 52 the vehicle was constantly monitored using short range localization
399 (SRL). Table 2, below, presents the missions conducted by the ISE Explorer vehicle B05 in the
400 Arctic during the campaign in 2010.
401
402 <Table 2 goes here>
403
404 Table 3, below, presents the survival estimates for missions 51, 52 and 53, considering the
405 unmitigated, mitigated and Bayesian updated survival profiles. The probability of survival for
406 missions 52 and 53, given the implementation of a monitoring distance of 31 km, was
407 calculated using Eq. (2).
408

409 < Table 3 goes here>

410

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

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

419 < Table 4 goes here>

420

421 Considering no failure mitigation, for the case where a monitoring distance of 87 km is
422 implemented (missions 51 and 52), the probability of surviving missions 55 and 56 increases
423 by 17.7%. On the other hand, if we consider only failure correction and the Bayesian inference,
424 with no monitoring distance, the probability of survival increases by 26.2%. This shows the
425 importance of quantifying the impact of failure correction and of updating the risk profile
426 based on subsequent missions.

427

428

429

430

431

432

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 sub-systems 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 whilst 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 the 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, pre-test, post-test, vehicle over-board (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*

457 of the probability of successful transition can be updated using the results of field
458 deployments, thus allowing us to update the availability prediction of an AUV.

459 In our view, this study has highlighted some limitations in the expert judgment elicitation,
460 as the practice of eliciting a single value for the probability of mitigation presents a problem.
461 Experts tend to assign a probability of 0 if the event is very unlikely and a probability of 1 if it
462 is very likely. In the elicitation of the probability of loss given a failure, the elicitation process
463 encourages the experts to assign a probability distribution, such as a beta distribution. A
464 similar approach should be adopted for eliciting the probability of failure mitigation. By doing
465 so, it allows the decision maker to update the risk for those failures in the light of subsequent
466 missions. This research shows that the practice of eliciting a single probability value for the
467 probability of failure mitigation can be problematic. It forces the analyst to make assumptions
468 regarding modelling the probability of mitigation. In making these assumptions the analyst
469 may introduce bias. This is evident with Figure 3; here the probability of failure mitigation is
470 represented as the number of counts. The probability of failure mitigation provided by
471 experts, were discrete values 0, 0.1, in increments of 0.1. This allowed us to model the
472 probability of failure mitigation as a number of counts. If a more detailed probability of failure
473 mitigation had been obtained, instead of using the number of counts to represent the
474 probability of failure mitigation could have used as density function. In the case study
475 presented in this paper, we modelled the probability of failure mitigation with a beta
476 distribution. In order to match a beta distribution to the probability of failure mitigation we
477 assumed that the mean of the beta distribution was the probability of failure mitigation. The
478 width of the prior (expert) probability distributions are also not known. To model the width,
479 we assumed that the prior variance was the same as the variance for the mode of the priori
480 probability of failure mitigation distribution, this was calculated using Eq. (7). Alternatively,

481 we could have assumed that the mode of the beta distribution was the probability of failure
482 mitigation or have made another assumption for the variance. This is, in our view, one
483 limitation of this research. This highlights the need for analysts to elicit a probability
484 distribution for the probability of failure mitigation. To our knowledge, this is not current
485 practice in risk modelling.

486 **Acknowledgments**

487 The authors would like to thank the anonymous reviewers for their very insightful comments.
488 We would also like to thank the experts who took part in the expert judgment elicitation
489 conducted in Halifax and in Vancouver.

490

APPENDIX A

Derivation of the Bayesian equation for three variables

Let Θ be a vector of probabilities of success for several independent binomial trials. Let n be a vector of the number of trials and let m be a vector of the number of successes.

The joint distribution P(Θ,m,n) can be calculated using the following equation

$$P(\theta, m, n) = P(\theta|m, n) \times P(m, n)$$

$$= P(\theta|m, n) \times P(m|n) \times P(n)$$

$$P(\theta, m, n) = P(m|\theta, n) \times P(\theta, n)$$

$$= P(m|\theta, n) \times P(\theta|n) \times P(n)$$

$$= P(m|\theta, n) \times P(\theta) \times P(n)$$

Θ is independent from n because we must know both m and n in order to infer Θ. Thus P(θ|n) = P(θ).

Taking into account the two terms for $P(\theta, m, n)$.

$$P(\theta|m, n) \times P(m|n) \times P(n) = P(m|\theta, n) \times P(\theta) \times P(n)$$

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

508 **References**

509 Brito, M. P., G. Griffiths, and P. Challenor, 2010: Risk Analysis for Autonomous Underwater
510 Vehicle Operations in Extreme Environments. *Risk Anal.*, **30**, 1771-1788.

511 Brito, M., G. Griffiths, 2011: A Markov Chain state transition approach to establishing critical
512 phases for AUV reliability. *IEEE J. Oceanic Eng.*, **36**, 139-149.

513 Brito, M., G. Griffiths, J. Ferguson, D. Hopkin, R. Mills, R. Pederson, and E. MacNeil, 2012: A
514 Behavioral Probabilistic Risk Assessment Framework for Managing Autonomous
515 Underwater Vehicle Deployments. *J. Atmos. Oceanic Technol.,* **29,** 1689-1703.

516 Brito, M. P., D. A. Smeed, and G. Griffiths, 2014: Underwater glider reliability and implications
517 for survey design. *J. Atmos. Oceanic Technol*., **31**, 2858-2870.

518 BS 4778, 1991. Quality vocabulary. Availability, reliability and maintainability terms. Guide to
519 concepts and related definitions. British Standards, London.

520 Crees, T., C. Kaminski, J. Ferguson, J. M. Laframboise, A. Forrest, J. Williams, E MacNeil, D.
521 Hopkin, and R. Pederson, 2010: UNCLOS under ice survey - An historic AUV
522 deployment in the Canadian high arctic. *Proc. OCEANS 2010 MTS/IEEE, 2010*, Seattle,
523 WA, 1-8.

524 Eriksen, C. C., T. J. Osse, R. D. Light, T. Wen, T. W. Lehman, P. L. Sabin, J. W. Ballard, and A. M.
525 Chiodi, 2001: Seaglider: A Long-Range Autonomous Underwater Vehicle for
526 Oceanographic Research. *IEEE J. Oceanic Eng.,* **26,** 424-436.

527 Feather, M. S., S. L. Cornford, 2003: Quantitative risk-based requirements reasoning. *Requir.*
528 *Eng.*, **8,** 248-265.

529 Griffiths, G., N. W. Millard, S. D. McPhail, P. Stevenson, and P. G. Challenor, 2003: On the
530 reliability of the Autosub autonomous underwater vehicle. *Underwater Technology*,
531 **25**, 175-184.

532 Griffiths, G., A. Trembanis, 2007: Eliciting expert judgment for the probability of AUV loss in
533 contrasting operational environments. *Proc. 15th International Symposium on Unmanned*
534 *Untethered Submersible Technology 2007*, Lee, NH, 1-17.
535 Griffiths, G., M. Brito, I. Robbins, M. Moline, 2009: Reliability of two REMUS-100 AUVs based
536 on fault log analysis and elicited expert judgment. *Proc. Autonomous Undersea*
537 *Systems Institute (AUSI) Int. Symp. on Unmanned Untethered Submersible Technology*
538 *2009*, Durham, NH, 1-12.
539 Kalbfleisch, J. D., and R. L. Prentice, 2002: *The statistical analysis of failure time data*. Wiley,
540 489 pp.
541 Kaplan, E. L., and P. Meier, 1958: Nonparametric estimation from incomplete observations. *J.*
542 *Am. Stat. Assoc.*, **53**, 457-481.
543 Kaminski, C., T. Crees, J. Ferguson, A. Forrest, J. Williams, D. Hopkin, and G. Heard, 2010: 12
544 days under ice – an historic AUV deployment in the Canadian High Arctic. *Proc.*
545 *IEEE/OES Autonomous Underwater Vehicles 2010*, Monterey, CA, pp. 1-11.
546 O'Hagan, A., C. E. Buck, A. Daneshkhah, J. R. Eiser, P. H. Garthwaite, D. J. Jenkinson, J. E.
547 Oakley, and T. Rakow, 2006: *Uncertain judgments: Eliciting experts' probabilities*.
548 Wiley, 338 pp.
549 Podder, T. K., M. Sibenac, H. Thomas, W. Kirkwood, and J. G. Bellingham, 2004: Reliability
550 growth of autonomous underwater vehicle-Dorado. *Proc. MTS/IEEE Conf. Oceans,*
551 *2004*, Kobe, Japan, 856-862.
552 Reason, J., 1990: *Human Error*. Cambridge University Press, 302 pp.
553 Rudnick, D. L., R. Davis, and J. T. Sherman, 2016: Spray Underwater Glider Operations. *J.*
554 *Atmos. Oceanic Technol.,* **6,** 1113-1122.

555 Singh, H., A. Can, R. Eustice, S. Lerner, N. McPhee, O. Pizarro, and C. Roman, 2004: Seabed
556 AUV Offers New Platform for High-Resolution Imaging. *Trans., Am. Geophys. Union,*
557 **85,** 294-295.
558 Subramanian, S., Elliott, L., Visnuvajjala, R. V., Tsai, W. T., and Mojdehbakhsh, R., 1996. Fault
559 mitigation in safety-critical software systems. *Proc. Ninth IEEE Symposium on*
560 *Computer-Based Medical Systems*, Ann Arbor, Michigan, pp.12-17.
561 Webb, D. C., P. J. Simonetti, and C. P. Jones, 2001: SLOCUM: an underwater glider propelled
562 by environmental energy. *IEEE J. Oceanic Eng.,* **26,** 447-452.

List of Tables

Table 1. Probability of failure mitigation for all faults presented in Brito et al (2012). The first column is the fault reference number. The second column contains the values of P_{Mi} for each failure. The third column P(loss|fault – 95% quantile) represents the probability of loss given the fault, without considering the mitigation. In the fourth column is the number of times that the fault has re-occurred. The fifth and sixth column present the hyper parameters of the beta distribution fitted to prior P_{Mi}. In columns seven to ten present the properties of the posterior PM. These columns present the hyper parameters, the mean and the standard deviation (sd). Column eleven presents the probability of loss given the fault and the mitigation, taking into account the prior P_{Mi}. Column twelve presents the probability of loss given the fault and the mitigation, considering the posterior P_{Mi}.

Fault ref. number	P_M	P(loss|Fault) - 95% quantile	Re-occurred? – number of times	Prior		Posterior				With mitigation 95% quantile	Bayesian updated 95% quantile
				Prior a	Prior b	a	b	mean	variance		
1	0	0.000413	1	-	-	-	-	-	-	4.13E-04	0.000413
2	0.9	6.4E-07	0	12.451	1.383	24.452	1.384	0.946	0.00189	6.40E-08	3.456E-08
3	0.8	0.0536	0	20.299	5.075	32.298	5.075	0.864	0.00306	1.07E-02	0.0072896
4	0.9	6.4E-07	0	12.451	1.383	24.452	1.384	0.946	0.00189	6.40E-08	3.456E-08
5	0.95	6.4E-07	0	6.489	0.341	18.489	0.341	0.982	0.000898	3.20E-08	1.152E-08
6	1	6.4E-07	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	0.00189	4.80E-02	0.02592
9	0.1	0.00464	0	3.575	32.175	15.575	32.175	0.326	0.00451	4.18E-03	0.0031274
10	0.95	1	0	6.489	0.341	18.489	0.341	0.982	0.000898	5.00E-02	0.018

Fault ref. number	P_M	P(loss|Fault) - 95% quantile	Re-occurred? – number of times	Prior		Posterior				With mitigation 95% quantile	Bayesian updated 95% quantile
				Prior a	Prior b	a	b	mean	variance		
11	0.95	6.4E-07	0	6.489	0.341	18.489	0.341	0.982	0.000898	3.20E-08	1.152E-08
12	1	0.504	0	-	-	-	-	-	-	0	0
13	0.8	0.0361	0	20.299	5.075	32.298	5.075	0.864	0.00306	7.22E-03	0.0049096
14	0.75	0.921	0	22.43	7.477	34.43	7.477	0.823	0.00342	2.30E-01	0.163017
15	0.4	0.901	0	45.493	68.239	57.492	68.239	0.457	0.00196	5.41E-01	0.489243
16	0.95	0.396	0	6.489	0.341	18.489	0.341	0.982	0.000898	1.98E-02	0.007128
17	1	0	0	-	-	-	-	-	-	0	0
18	0.9	0.0166	0	12.451	1.383	24.452	1.384	0.946	0.00189	1.66E-03	0.0008964
19	0.8	0	0	20.299	5.075	32.298	5.075	0.864	0.00306	0.00E+00	0
20	0.4	0.79	0	45.493	68.239	57.492	68.239	0.457	0.00196	4.74E-01	0.42897
21	0.8	0.0361	0	20.299	5.075	32.298	5.075	0.864	0.00306	7.22E-03	0.0049096
22	0	1	0	-	-	-	-	-	-	1	1
23	0.9	0	0	12.451	1.383	24.452	1.384	0.946	0.00189	0.00E+00	0
26	0.95	0.0361	0	6.489	0.341	18.489	0.341	0.982	0.000898	1.81E-03	0.0006498
28	0.95	0.167	0	6.489	0.341	18.489	0.341	0.982	0.000898	8.35E-03	0.003006
29	0.9	0	0	12.451	1.383	24.452	1.384	0.946	0.00189	0.00E+00	0

1

2

Fault ref. number	P_M	P(loss|Fault) - 95% quantile	Re-occurred? – number of times	Prior		Posterior				With mitigation 95% quantile	Bayesian updated 95% quantile
				Prior a	Prior b	a	b	mean	variance		
30	0.8	0.176	0	20.299	5.075	32.298	5.075	0.864	0.00306	3.52E-02	0.023936
31	0.95	0.0361	0	6.489	0.341	18.489	0.341	0.982	0.000898	1.81E-03	0.0006498
32	1	0	0	-	-	-	-	-	-	0	0
33	1	0	0	-	-	-	-	-	-	0	0
34	0.9	0.00983	0	12.451	1.383	24.452	1.384	0.946	0.00189	9.83E-04	0.0005308
35a	1	0	0	-	-		-	-	-	0	0
35b	0.6	1	0	68.239	45.493	80.239	45.493	0.638	0.00182	4.00E-01	0.362
36	0.9	0.798	0	12.451	1.383	24.452	1.384	0.946	0.00189	7.98E-02	0.043092
37	0.1	0.00464	0	3.575	32.175	15.575	32.175	0.326	0.00451	4.18E-03	0.0031274
38	1	0	0	-	-	-	-	-	-	0	0
39	0.75	0.109	0	22.43	7.477	34.43	7.477	0.823	0.00342	2.73E-02	0.019293
40	0.5	0.202	0	59.256	59.256	71.256	59.256	0.546	0.00189	1.01E-01	0.091708
41	0.5	0.0189	0	59.256	59.256	71.256	59.256	0.546	0.00189	9.45E-03	0.0085806
42	0.5	0.0197	2	59.256	59.256	69.256	61.256	0.531	0.00189	9.85E-03	0.0092393
43	0.5	1	0	59.256	59.256	71.256	59.256	0.546	0.00189	5.00E-01	0.454
44	0.5	0.0197	2	59.256	59.256	69.256	61.256	0.531	0.00189	9.85E-03	0.0092393

1

1

2

Fault ref. number	P_M	P(loss|Fault) - 95% quantile	Re-occurred? – number of times	Prior		Posterior				With mitigation 95% quantile	Bayesian updated 95% quantile
				Prior a	Prior b	a	b	mean	variance		
45	0	0.0191	0	-	-	-	-	-	-		
										0.0192	0.0191
46	0.5	0.00009	2	59.256	59.256	69.256	61.256	0.531	0.00189	4.50E-05	4.221E-05
47	0.5	0.78	0	59.256	59.256	71.256	59.256	0.546	0.00189	3.90E-01	0.35412
48	0.5	0.451	0	59.256	59.256	71.256	59.256	0.546	0.00189	2.26E-01	0.204754
50	0.5	1	0	59.256	59.256	71.256	59.256	0.546	0.00189	5.00E-01	0.454
51	0.95	0.0361	0	6.489	0.341	18.489	0.341	0.982	0.000898	1.81E-03	0.0006498
52	0.5	0.78	0	59.256	59.256	71.256	59.256	0.546	0.00189	3.90E-01	0.35412
53	0.5	0.78	0	59.256	59.256	71.256	59.256	0.546	0.00189	3.90E-01	0.35412
54	0.1	0.00947	1	3.575	32.175	14.575	33.175	0.305	0.00435	8.52E-03	0.0065817

3

4

Table 2. All mission conducted by ISE Explorer B05 in the Arctic Campaign of 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, 52 and 53. The monitoring mission of 31km corresponds to mission 51.

	Probability of survival			
Scenario	Mission 51	Mission 52	Mission 53	Survival for all missions
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 31km		0.899	0.894	0.8048
Bayesian updated + monitor 31km		0.905	0.905	0.8201

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 87km	0.610	0.610	0.372
Mitigated+monitor 87km	0.802	0.802	0.644
Bayesian updated+monitor 87km	0.817	0.817	0.667

List of figures captions

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.

Fig.3: Probability of failure mitigation. In gray, the prior of the probability of mitigation distribution. In black, the posterior probability of mitigation distribution.

Fig. 4: Survival distribution for ISE Explorer vehicle deployment in the Arctic. In full line is the survivability without mitigation. In dashed line is the survivability considering the mitigation based on the prior belief. In dotted line is the survivability considering the mitigation updated.

List of Figures

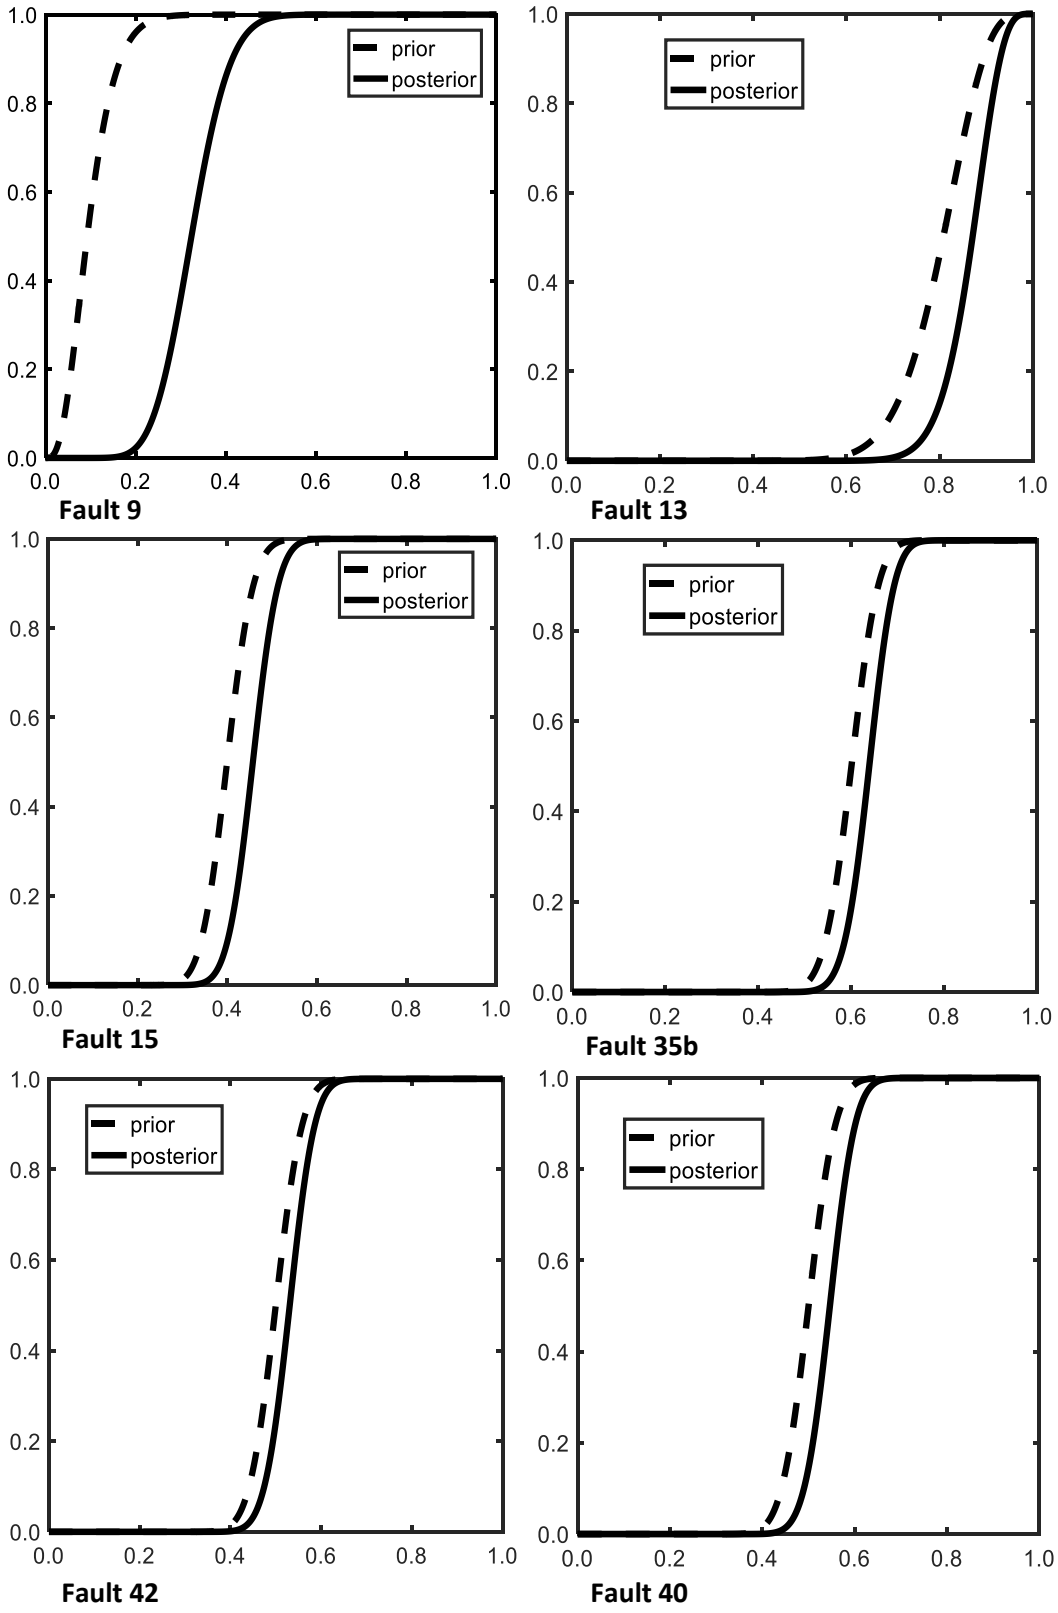

Fig. 1. Density distribution for the mean of the probability of fault mitigation for failures 9, 13, 15, 35b, 42 and 40.

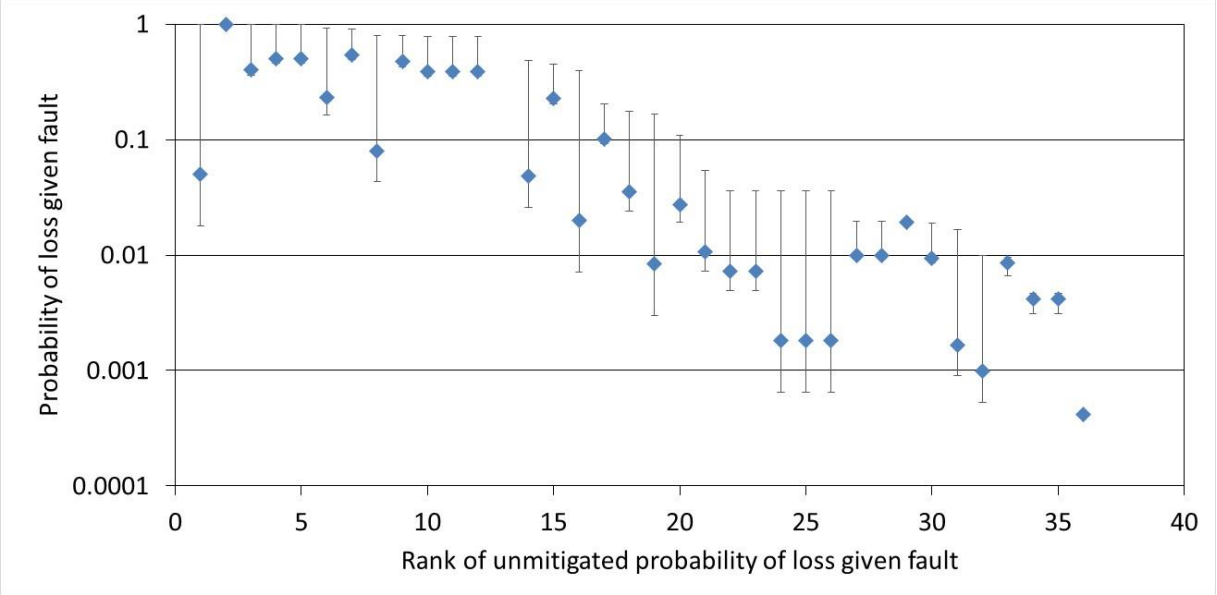

Fig. 2. Probability of loss given fault, mitigation and trials results. Ranked order by unmitigated probability of loss given fault, for those faults above 0.01. The upper point of the "error bar" is the unmitigated data, the main point is after a priori estimated mitigation and the lower error bar is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

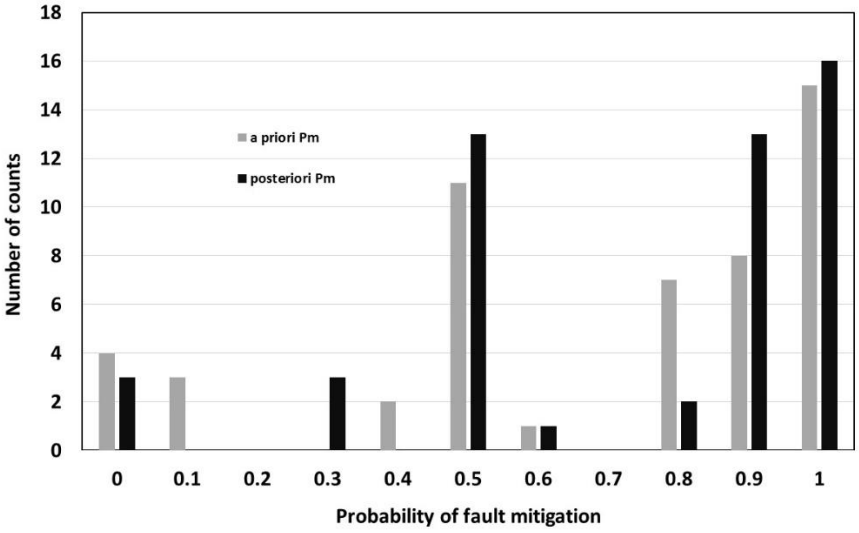

Fig. 3. Probability of failure mitigation. In gray, the prior of the probability of mitigation distribution. In black, the posterior probability of mitigation distribution.

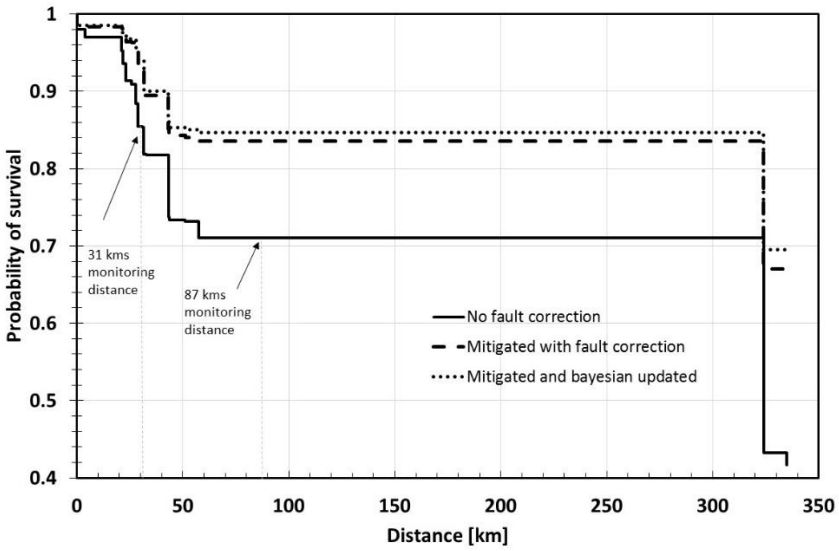

Fig. 4. Survival distribution for ISE Explorer vehicle deployment in the Arctic. In full line is the survivability without mitigation. In dashed line is the survivability considering the mitigation based on the prior belief. In dotted line is the survivability considering the mitigation updated using Bayesian updated.