European Journal of Operational Research 255 (2016) 259–271

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Allocation of tasks for reliability growth using multi-attribute utility

Kevin J. Wilson[a,*], John Quigley[b,1]

[a] School of Mathematics and Statistics, Newcastle University, UK
[b] Department of Management Science, University of Strathclyde, UK

ARTICLE INFO

Article history:
Received 8 January 2015
Accepted 7 May 2016
Available online 27 May 2016

Keywords:
Utility theory
Reliability growth
Bayesian experimental design
Multivariate risk aversion
Expert judgement

ABSTRACT

In reliability growth models in particular, and project risk management more generally, improving the reliability of a system or product is limited by constraints on cost and time. There are many possible tasks which can be carried out to identify and design out weaknesses in the system under development. This paper considers the allocation problem: which subset of tasks to undertake. While the method is applicable to project risk management generally, the work has been motivated by reliability growth programmes. We utilise a model for reliability growth, based on an efficacy matrix, developed with engineering experts in the aerospace industry. We develop a general multi-attribute utility function based on targets for cost, time on test and system reliability. The optimal subset is identified by maximising the prior expected utility. We derive conditions on the model parameters for risk aversion and loss aversion based on observed properties of preference. We give conditions for multivariate risk aversion under the general form of the utility function. The method is illustrated using an example informed by work with aerospace organisations.

1. Introduction

Selecting a programme of activities optimal against multiple criteria is cognitively challenging and time consuming for decision makers but can be aided with appropriate decision support tools if preferences can be represented mathematically. In the development of large, complex products or systems, the system is analysed at various stages for potential design weaknesses and, once weaknesses have been identified, they are designed out. This improves the system's reliability. Examples of tasks which are used to identify weaknesses are fault tree analysis, failure modes and effects analysis, test, analyse and fix (TAAF), load strength analysis, vibration testing, simulation studies and accelerated life testing (O'Connor, 1991).

The outcomes of these tasks will not be mutually exclusive: tasks may expose multiple weaknesses and weaknesses may be exposed by various tasks. There will typically be neither the budget nor the time to carry out all of the potential reliability tasks. Therefore engineers choose and sequence a subset of tasks to improve the product's reliability. This paper considers methods to select such a portfolio of reliability tasks.

While our motivation is concerned with managing reliability growth programmes, trading between performance targets, project duration and costs dynamically throughout a project is a concern for project risk management broadly. Previous approaches have used mathematical optimisation (Goel, Grievink, & Weijnen, 2003; Gurov, Utkin, & Shubinsky, 1995; Hsieh, 2003; Hsieh & Hsieh, 2003) or fuzzy logic (Idrus, Nuruddin, & Rohman, 2011; Nieto-Morote & Ruz-Vila, 2011; Petrović, Petrović, & Vujošević, 1996; Petrovic & Aköz, 2008) to solve the decision problem. We use a utility-based approach. If we have hard constraints, like in optimisation, we can miss some desirable solutions. As such, we need to develop methods that penalise as we move farther away from desirable targets. We have chosen to develop the decision support on utility theory, as we seek to represent preference trade-offs rather than vagueness of decision makers. Multi-attribute utility has been used in a similar manner in the area of portfolio resource allocation (Aouni, Colapinto, & Torre, 2014; Hallerbach, Ning, Soppe, & Spronk, 2004; Salo, Keisler, & Morton, 2011) and the simplifying assumption of utility independence is identified as desirable to specify a utility function.

In the context of reliability growth (Quigley & Walls, 2006; Johnson, Quigley, & Walls, 2006) developed a model which aimed to represent the process experienced by engineers. It explicitly considered all of the potential faults and tasks to identify them resulting in the use of an efficacy matrix. The efficacy of each task is assessed against each potential failure mode producing an efficacy matrix to measure the conditional probability of exposing

---

[*] Corresponding author. Tel.: +441912087311.
E-mail addresses: k.j.wilson1984@gmail.com, kevin.wilson@ncl.ac.uk (K.J. Wilson), j.quigley@strath.ac.uk (J. Quigley).
[1] Tel.: +44 141 548 3152.

http://dx.doi.org/10.1016/j.ejor.2016.05.014

the failure mode given its presence in the design. Such a matrix could have uses across project risk management problems. Reliability improves as specific design weaknesses are identified and removed from the system. All of the parameters in the model can be elicited from observable quantities. We use this model as the basis to solve the decision problem of task allocation. Johnson et al. (2006) also considered task allocation and outlined an integer programming approach which minimised costs subject to constraints on expected reliability and time on test. The shortcoming of this approach is that it provides little sensitivity around the reliability and time targets: an allocation which just failed to meet the targets was unacceptable and any allocation which met the targets was equally desirable.

In this paper we propose a Bayesian solution to the task allocation problem; choosing the allocation which maximises the prior expectation of a utility function representing the engineers' preferences over cost, reliability and time on test. A general utility function over these attributes, which utilises the idea of mutually utility independent hierarchies, will be developed. The form of this utility function will be adapted to satisfy observed properties of marginal preferences from decision makers in experiments. In particular, we develop conditional utility functions to represent risk averse preferences and loss averse preferences which satisfy the isolation effect. That is, preferences of individuals over lotteries generally discard elements that the lotteries have in common (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992).

We consider the impact of the form of the utility function on preferences over multiple attributes and give conditions for the individual risk averse and loss averse utilities to lead to multivariate risk aversion (Richard, 1975). The resulting optimal allocation is more sensitive to small changes in expected reliability and time on test around the target values than the integer programming approach. This is the first time multi-attribute utility has been used for task allocation in reliability growth modelling.

The contribution of the paper takes two forms; a theoretical contribution on multi-attribute utility and a methodological contribution on reliability growth specifically and project risk management more generally. In the first case, we consider for the first time the implications for multivariate preference behaviour by assuming utility independence within a mutually utility independent hierarchy (MUIH). Proposition 4 shows that such structures are sufficiently flexible to represent multivariate risk aversion, risk neutrality and risk seeking behaviour. Proposition 3 shows that not all attributes within a MUIH are by necessity utility independent. The illustrative example quantifies the impact of assuming different preference behaviours of the decision maker within utility functions. The preferences of the decision maker over multiple attributes can result in different optimal allocations of reliability tasks. In the second case, we develop a methodology within a reliability growth framework which allows engineers to make decisions about which activities to undertake which explicitly considers trade-offs between the important attributes in their decision. The methodology captures varying preference behaviours and gives an analytically tractable solution to the decision problem. We indicate the generalisation of the methodology to similar decision problems in project risk management.

The rest of the paper is structured as follows. In Section 2 we outline the model for reliability growth developed by Johnson et al. (2006). In Section 3 we outline our Bayesian expected utility approach to the allocation problem, giving the general form of the solution method and developing utility functions over reliability, cost and time on test. In Section 4 we present an illustrative example to compare the expected utility approach to the integer programming approach and to investigate the properties of the expected utility approach. We provide a simulation study to investigate the effects of assuming risk aversion and loss aversion of the decision maker. Finally, we summarise the paper and identify future work in Section 5.

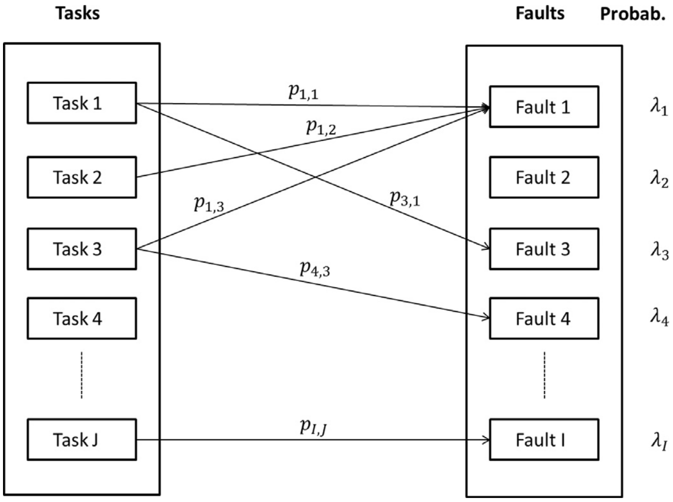

Fig. 1. A diagram illustrating the form of the efficacy matrix.

2. An expert judgement informed reliability growth model

We adopt the approach developed in Quigley and Walls (2006) and Johnson et al. (2006). In Section 2.1 we define the efficacy matrix, which is core to the reliability growth model. In Section 2.2 we derive the reliability assessment for a design prior to undertaking reliability tasks. In Section 2.3 we derive the updated reliability assessment following the outcome of a reliability task.

2.1. The efficacy matrix

Suppose that the current design of an engineering system has associated with it a number of identified potential faults, labelled $i = 1, \ldots, I$. Then, for each fault i, there is some probability, denoted λ_i, that this fault will be realized as a failure at some point in the lifetime of the system. Define X_i to be an indicator variable,

$$X_i = \begin{cases} 1, & \text{if fault } i \text{ is ever realised,} \\ 0, & \text{otherwise.} \end{cases}$$

The probabilities of X_i being in its two possible states are λ_i and $1 - \lambda_i$ respectively.

As part of the growth programme there are a number of possible tasks which could be performed on the system, labelled $j = 1, \ldots, J$. Each task will have a certain efficacy at identifying each of the faults in the system. Denote by $p_{i,j}$ the conditional probability that task j will realise fault i given that the fault exists within the system.

An illustration of the efficacy matrix is given in Fig. 1. We see the J possible tasks to identify the I potential faults in the system. Each of the faults has an associated probability that it exists in the system. In the figure, task 1 will identify faults 1 and 3 with probabilities $p_{3,1}$, $p_{1,1}$ respectively. Fault 1 could also be identified by tasks 2 and 3. Therefore there are multiple routes which could identify fault 1. By contrast, none of the tasks in the figure can identify fault 2. The probability that fault 2 exists, λ_2, will never change as a result of performing any task.

We can elicit both λ_i and $p_{i,j}$, by asking questions about observable quantities, from engineering experts inside the organisation. For more information see Hodge, Evans, Quigley, and Walls, 2001. Similarly, Walls and Quigley (1999) developed a Bayesian model based on observable quantities for reliability growth in the TAAF cycle.