Optimal expert elicitation to reduce interval uncertainty

Nadia Ben Abdallah, Sébastien Destercke

Optimal expert elicitation to reduce interval uncertainty

Nadia Ben Abdallah
Heudiasyc laboratory
University of Technology of Compiègne
Compiègne, France
nadia.ben-abdallah@hds.utc.fr

Sébastien Destercke
Heudiasyc laboratory
University of Technology of Compiègne
Compiègne, France
sebastien.destercke@hds.utc.fr

Abstract

Reducing uncertainty is an important problem in many applications such as risk and reliability analysis, system design, etc. In this paper, we study the problem of optimally querying experts to reduce interval uncertainty. Surprisingly, this problem has received little attention in the past, while similar issues in preference elicitation or social choice theory have witnessed a rising interest. We propose and discuss some solutions to determine optimal questions in a myopic way (one-at-a-time), and study the computational aspects of these solutions both in general and for some specific functions of practical interest. Finally, we illustrate the application of the approach in reliability analysis problems.

1 INTRODUCTION

When data on some quantity or model of interest is sparse or non-existing, elicitation, i.e., the process of extracting human judgement through questions, is often a valuable and sometimes the unique source of additional knowledge. There is a substantial literature dating back to the sixties on elicitation and is mainly related to probability encoding (Winkler, 1969; Spetzler and Stael von Holstein, 1975) and preference elicitation (Keeney et al., 1979). Elicitation is used in a broad range of fields including risk assessment (Cooke, 1991), reliability analysis, preference model elicitation (Viappiani and Kroer, 2013; Guerin et al., 2013), etc. to support assessment and decision making.

A critical part of the elicitation is then how to choose the questions to ask. Those need to be simple (i.e., do not require high cognitive effort) and in terms and format experts are familiar with. Furthermore, when the elicitation is conducted to reach some objective, for instance bringing an answer to a question, selecting the best alternative in a set, or estimating some quantity with a desired level precision, the process of information acquisition need to be optimal for the elicitation to be effective and the least possible time or effort consuming.

How to choose sequences of optimal questions, or even the notion of optimal queries, has received surprisingly little attention when the aim is to reduce our uncertainty over some quantities. Indeed, the great majority of techniques to do so prescribe generic questions, without considering the consequences of answers on some final goal (Aspinall and Cooke, 2013) (the work of Curtis and Wood (Curtis and Wood, 2004), settled in a probabilistic context, is an exception). This contrasts with other fields such as preference elicitation of social choice theory, with works dating back two decades ago (Boutilier et al., 1997; Wang and Boutilier, 2003; Boutilier et al., 2006) and still thriving today (Viappiani and Kroer, 2013; Benabbou et al., 2014; Boutilier et al., 2013).

The goal of this paper is to explore similar ideas when the goal is to reduce interval uncertainty by asking successive simple questions to the experts. We want to develop querying strategies that are adaptive and optimal, i.e., that select at each stage of the elicitation the best questions based on the answers to the previous ones. In this paper, we focus on so-called myopic (Wang and Boutilier, 2003; Chajewska and Koller, 2000) strategies, where optimal questions are selected one-at-a-time.

The remainder of the paper is organized as follows. In Section 2, we formalize the sequential elicitation model for the problem of interval uncertainty reduction in the general case. Within this same section (Section 2.3), we describe different query selection strategies, and analyse their computational costs in the general case, which is an important aspect to consider in adaptive procedures. Section 3 then discusses the case of specific yet important (in practice) type of functions, namely monotonic and multi linear functions. In the last section, we illustrate how the approach can be used in reliability analysis.

2 GENERAL FRAMEWORK

2.1 PROBLEM STATEMENT

Let Φ be a function mapping a set of n logically independent inputs $(x_1, ..., x_n)$, each of them being defined on X_i, to an output y in Y :

$$\begin{aligned} \Phi \quad : \quad \mathbf{X} = \times_{i=1...n} X_i \quad &\rightarrow \quad Y \\ \mathbf{x} = (x_1, ..., x_n) \quad &\mapsto \quad \Phi(\mathbf{x}) = y. \end{aligned}$$

In this paper, we are interested in the situation where x_i is a precise but ill-known value, whose uncertainty is described by an interval $X_i = \left[\underline{X}_i, \overline{X}_i\right] \subset \mathbb{R}$ of the real line. Such kind of uncertainty, where the true value is exact, is sometimes called epistemic (by opposition to aleatory). A natural way to quantify the amount of uncertainty in X_i is by its width

$$U_{X_i}(x_i) = U_{X_i} = \overline{X_i} - \underline{X_i}.$$

We also require the function Φ to be continuous, so that the response y corresponding to the initial state of knowledge on the inputs lies in the bounded interval :

$$Y = \Phi(\mathbf{X}) = [\min_{\mathbf{x}\in\mathbf{X}} \Phi(\mathbf{x}), \max_{\mathbf{x}\in\mathbf{X}} \Phi(\mathbf{x})] = [\underline{Y}, \overline{Y}]. \quad (1)$$

Example 1. *Consider the function* $\Phi(x_1, x_2, x_3) = x_1x_2 - x_2x_3$ *with* $X_1 = X_2 = X_3 = [0, 1]$*, then we have*

$$\underline{Y} = \Phi(\underline{X}_1, \overline{X}_2, \overline{X}_3) = -1; \overline{Y} = \Phi(\overline{X}_1, \overline{X}_2, \underline{X}_3) = 1.$$

The problem we are considering is the following : we want to reduce our uncertainty $U_Y = \overline{Y} - \underline{Y}$ by asking question to experts, to attain some objectives. For instance, we may want to reduce the uncertainty under some threshold $U_Y \leq s_0$ or simply reduce the most U_Y in a given number of questions. As expert elicitation is time-consuming and cognitively demanding for the expert, and economically expensive for the decision maker, we want to ask as few questions as possible, or to be the most effective possible on those questions we ask. In other words, we want the querying strategy to be optimal. This is what we develop in the next sections.

2.2 QUERIES AND ANSWERS

In expert elicitation in general, and when the elicitation is made of many successive questions, it is important to use simple questions that not require high cognitive effort (for understanding and answering) for the expert to be efficient throughout the interview. Possible simple queries formats include local bound queries ("$x_i \leq \alpha$?"), pairwise comparison judgements ("$x_i \leq x_j$" ?), etc. (Braziunas and Boutilier, 2007).

In our method, we use questions of the type "$x_i \leq \alpha$?", with $\alpha \in X_i$. We denote such a query Q_i^α and the set of possible queries $\mathcal{Q} = \{Q_i^\alpha, i \in N = \{1, 2, ..., n\}, \alpha \in X_i\}$.

In the particular case of local bound queries, the set of possible answers $\mathcal{A}$ is binary : $\mathcal{A} = \{Yes, No\}$. We recall that, for simplicity and for conciseness, we assume that the expert is an oracle, so the "I don't know" answer is not considered here [1]. Note that the ideas presented in the paper could easily be applied to other sets of questions/answers $\mathcal{Q}, \mathcal{A}$, yet binary questions are the simplest and the most natural to ask to experts.

When a question Q_i^α is asked and answer $A \in \mathcal{A}$ is given, X_j remains unchanged for every $j \neq i$, while X_i is updated to $X_i(Q_i^\alpha, A)$ as follows :

$$X_i(Q_i^\alpha, A) = \begin{cases} X_i \cap [-\infty, \alpha] & \text{if } A = Yes \\ \\ X_i \cap [\alpha, -\infty] & \text{if } A = No \end{cases} \quad (2)$$

which satisfies $X_i(Q_i^\alpha, A) \subseteq X_i$ and $X_j(Q_i^\alpha, A) = X_j$ for every $j \neq i$. Consequently, the output uncertainty set is updated from Y into $Y(Q_i^\alpha, A)$:

$$Y(Q_i^\alpha, A) = \Phi\left(X_{-i} \times X_i(Q_i^\alpha, A)\right), \quad (3)$$

where $X_{-i} = \times_{j\neq i} X_j$ denotes the Cartesian product of all unchanged intervals. As for any $Q \in \mathcal{Q}$ and $A \in \mathcal{A}$ we have $Y(Q, A) \subseteq Y$ by simple interval inclusion, the following relation always holds :

$$U_Y \geq U_{Y(Q,A)} \quad (4)$$

therefore ensuring an uncertainty reduction.

Example 2. *In Example 1, assume we ask the question* $Q_1^{0.5}$ *and receive the answer* Yes*, then*

$$X_1(Q_1^{0.5}, Yes) = [0, 0.5]$$

$$\overline{Y}(Q_1^{0.5}, Yes) = \Phi(\overline{X}_1, \overline{X}_2, \underline{X}_3) = 0.5.$$

2.3 QUERY SELECTION STRATEGIES

A query selection strategy corresponds to define and choose optimal questions. There are two main ways to do so : myopically, where questions are selected and asked one at a time, successively, and sequentially, where the set of successive questions is selected globally. Here, we retain the myopic approach for the following reasons : it is often simpler to solve, sometimes allowing for analytical exact solutions, and does not require to specify the number of asked questions in advance, a particularly interesting feature in iterative and interactive querying process.

1. Should the expert return "I don't know" to Q_i^α, then a simple strategy is to remove Q_i^α (and possibly questions with similar values of α) from the question set $\mathcal{Q}$ and then select the optimal one among the remaining ones.

The selection process of the myopic approach consists in solving the following optimization problem at each iteration :

$$Q^* = \arg\min_{Q \in \mathcal{Q}} U_{Y(Q)}, \quad (5)$$

where $U_{Y(Q)}$ is the uncertainty reduction induced by query Q. However, as the answer $\mathcal{A}$ that will be given to Q is unknown, we face a typical problem of decision making under uncertainty.

In our case, the decision is a couple $(i, \alpha) \in N \times X_i$, the uncertain event is the answer to the question, and the outcome we want to maximize is the uncertainty reduction in the output Y. Re-writing the decision problem using notations of our query selection problem leads to the following characterization of the optimal queries :

$$Q^* = (i^*, \alpha^*) = \arg\min_{i \in N} \min_{\alpha \in X_i} U_{Y(Q_i^\alpha)}, \quad (6)$$

which is a two stage optimization problem. First, we determine the optimal local bound value for each input i, and calculate the uncertainty reduction induced by that local query. Then, we select the entity i^* that leads to the highest uncertainty reduction in y.

Algorithm : Iterative elicitation for uncertainty reduction

$Inputs : X_i (i \in N), s_0$
while $U_Y >= s_0$ **do**
  for i in N **do**
    Compute $\alpha^* = \arg\min_{\alpha \in X_i} U_{Y(Q_i^\alpha)}$
    Compute $U_{Y(Q_i^{\alpha^*})}$
  end for
  $i^* = \arg\min_{i \in N} U_{Y(Q_i^{\alpha^*})}$
  Ask query : "$x_{i^*} <= \alpha^*$?"
  Obtain answer
  Update X_{i^*}
  Compute U_Y
end while

In the following, we describe the computations involved in the first optimization step (the computation of $U_{Y(Q_i^\alpha)}$) for a given $i \in N$ for different decision criteria.

2.3.1 Maximin strategy

The maximin strategy corresponds to a pessimistic view, where the value $U_{Y(Q_i^\alpha)}$ corresponds to the answer that yields the lowest uncertainty reduction :

$$\begin{cases} U^{Mm}_{Y(Q_i^\alpha)} = \max_{A \in \mathcal{A}} U_{Y(Q_i^\alpha, A)} \\ \alpha^{*,Mm} = \arg\min_{\alpha \in X_i} \max(U_{Y(Q_i^\alpha, No)}, U_{Y(Q_i^\alpha, Yes)}) \end{cases}$$

We can show that solving the optimization problem to get the optimizing α^* is equivalent to finding the intersection of the two functions $U_{Y(Q_i^\alpha, Yes)}, U_{Y(Q_i^\alpha, No)}$ of α. The following propositions indicates that a general and efficient method to find the solution is to use a dichotomy search on the space $[\underline{X}_i, \overline{X}_i]$

Proposition 1. *Functions $U_{Y(Q_i^\alpha, Yes)}$ and $U_{Y(Q_i^\alpha, No)}$ measuring the uncertainty level on Y induced by a positive and a negative answer to Q_i^α and defined on X_i are*
– *increasing and decreasing in α, respectively, and*
– *intersect at least once at $M_i \subset X_i$ (M_i is a single point or an interval).*

Proof. *We have that*

$$U_{Y(Q_i^\alpha, Yes)} = \max_{X_{-i} \times [\underline{X}_i, \alpha]} \Phi(\mathbf{x}) - \min_{X_{-i} \times [\underline{X}_i, \alpha]} \Phi(\mathbf{x}) \quad (7)$$

To show that $U_{Y(Q_i^\alpha, Yes)}$ is increasing in α, we need to show that $U_{Y(Q_i^\alpha, Yes)} \leq U_{Y(Q_i^\beta, Yes)}$ for $\alpha \leq \beta$. This result follows from $X_{-i} \times [\underline{X}_i, \alpha] \subseteq X_{-i} \times [\underline{X}_i, \beta]$.

The same reasoning can be applied to

$$U_{Y(Q_i^\alpha, No)} = \max_{X_{-i} \times [\alpha, \overline{X}_i]} \Phi(\mathbf{x}) - \min_{X_{-i} \times [\alpha, \overline{X}_i]} \Phi(\mathbf{x}) \quad (8)$$

to show that $U_{Y(Q_i^\alpha, No)}$ is decreasing in α.

To demonstrate the second part of the proposition, simply observe that :

$$\max_{\alpha \in X_i} U_{Y(Q_i^\alpha, Yes)} = U_{Y(Q_i^{\overline{X}_i}, Yes)} = U_Y,$$

$$\max_{\alpha \in X_i} U_{Y(Q_i^\alpha, No)} = U_{Y(Q_i^{\underline{X}_i}, No)} = U_Y,$$

where U_Y is the uncertainty before the question. As both functions have the same maximum, are continuous (since Φ is), and are respectively increasing and decreasing in α, they have at least one point of intersection.

When M_i is an interval $[\underline{M}_i, \overline{M}_i]$, we simply take the middle point $\alpha^{*,Mm} = \underline{M}_i + \overline{M}_i / 2$. In some situations, it may also happen that the intersection occurs on the bounds of X_i for all i, which means that the proposed optimal question is likely to be uninformative, unless the expert answer reduces the interval $[\underline{X}_i, \overline{X}_i]$ to a point, which is unlikely. When such a scenario occurs, we use a different strategy that defines optimality as the highest reduction of uncertainty, no more on Y, but on X_i. This heuristic is equivalent, when X_is are intervals, to choosing the largest interval $i^* = \arg\max_i U_{X_i}$ and to pick the mid of this interval, i.e., $\alpha^{*,Mm} = \frac{\underline{X}_{i^*} + \overline{X}_{i^*}}{2}$.

In Section 3, we will show that for specific functions, there are more efficient ways than a naive dichotomic search to determine $\alpha^{*,Mm}$.

2.3.2 Maximax strategy

While the maximin strategy is pessimistic, the maximax strategy is optimistic and takes as value $U_{Y(Q_i^\alpha)}$ the answer that yields the highest uncertainty reduction :

$$\begin{cases} U^{MM}_{Y(Q_i^\alpha)} = \min_{A\in\mathcal{A}} U_{Y(Q,A)} \\ \\ \alpha^{*,MM} = \arg\min_{\alpha\in X_i} \min(U_{Y(Q_i^\alpha,No)}, U_{Y(Q_i^\alpha,Yes)}). \end{cases}$$

It is straightforward from Proposition 1 that the local bound optimization step leads to an optimal value $\alpha^{*,MM}$ that coincides either with the upper or lower bound of X_i. The maximax strategy is therefore not interesting in our problem, as it will lead to questions that are most likely to receive a useless answer. We will therefore not retain this approach in this paper.

2.3.3 Hurwicz' strategy

Hurwicz's strategy evaluates the value of a question Q by a convex combination between the maximin and maximax strategies. It therefore allows to go from a pessimistic to an optimistic point of view and reads :

$$\begin{cases} U^{H(p)}_{Y(Q_i^\alpha)} = pU^{MM}_{Y(Q)} + (1-p)U^{Mm}_{Y(Q)} \\ \\ \alpha^{*,H(p)} = \arg\min_{\alpha\in X_i}(p\min(U_{Y(Q_i^\alpha,Yes)}, U_{Y(Q_i^\alpha,No)}) + \\ \qquad\qquad (1-p)\max(U_{Y(Q_i^\alpha,Yes)}, U_{Y(Q_i^\alpha,No)})). \end{cases}$$

Here, $p \in [0,1]$ is an optimism coefficient, and we retrieve the maximax and maximin strategies when $p = 1$ and $p = 0$, respectively. Note that we can use the fact that for any $\alpha \leq \underline{M}_i$ ($\alpha \geq \overline{M}_i$), we have $U_{Y(Q_i^\alpha,No)} >= U_{Y(Q_i^\alpha,Yes)}$ ($U_{Y(Q_i^\alpha,No)} <= U_{Y(Q_i^\alpha,Yes)}$) to rewrite the above equations into :

$$\begin{cases} L = \min_{\alpha\in[\underline{X}_i, \inf M_i]}(pU_{Y(Q_i^\alpha,Yes)} + (1-p)U_{Y(Q_i^\alpha,No)}) \\ R = \min_{\alpha\in[\sup M_i, \overline{X}_i]}(pU_{Y(Q_i^\alpha,No)} + (1-p)U_{Y(Q_i^\alpha,Yes)}) \\ \alpha^{*,H(p)} = \min(L,R). \end{cases}$$

2.3.4 Bayesian strategy

Up to now, we have not considered any a priori information about the likelihood of answering Yes or No. However, this can lead to consider very unlikely answers, such as answering Yes to $Q_i^{\underline{X}_i}$ (as is the case in the maximax strategy). One alternative is then the Bayesian strategy, where we assume the existence of a probability distribution P over the set of answers $\mathcal{A}$, this probability modelling our subjective beliefs about the likelihood of getting the different answers. We then evaluate a query Q by the expected reduction $\mathbb{E}_P(U_{Y(Q,A)})$ of uncertainty on y induced by the possible answers :

$$\begin{cases} U^{B}_{Y(Q_i^\alpha)} = \mathbb{E}_P(U_{Y(Q_i^\alpha,A)}) = \sum_{A\in\mathcal{A}} P(A)U_{Y(Q_i^\alpha,A)} \\ \\ \alpha^{*,B} = \arg\min_{\alpha\in X_i}(P(Yes|Q_i^\alpha)U_{Y(Q_i^\alpha,Yes)} + \\ \qquad\qquad P(No|Q_i^\alpha)U_{Y(Q_i^\alpha,No)})). \end{cases}$$

When the available evidence suggests that a quantity x_i lies in an interval X_i, it is common to follow Laplace's indifference principle and quantify uncertainty by assuming a uniform probability distribution over that set. Under this assumption, the probability of the positive and negative answers to a question Q_i^α are proportional to the width of the sub-interval of X_i they lead to :

$$P(Yes|Q_i^\alpha) = P(\underline{X}_i \leq x_i \leq \alpha) = \frac{\alpha - \underline{X}_i}{\overline{X}_i - \underline{X}_i}$$

and

$$P(No|Q_i^\alpha) = P(\alpha \leq x_i \leq \overline{X}_i) = \frac{\overline{X}_i - \alpha}{\overline{X}_i - \underline{X}_i}.$$

These probabilities can then be modified according to the information we have (for instance, if we have reasons to think that the true value is closer to $\overline{X}_i$).

Remark 1. *Note that when $U_{Y(Q_i^\alpha,Yes)} = U_{Y(Q_i^\alpha,No)} = U_{Y(Q_i^\alpha)}$, then $\mathbb{E}_P(U_{Y(Q_i^\alpha,A)}) = U_{Y(Q_i^\alpha)}$, whatever the values of P. This means, among other things, that the function $U^B_{Y(Q_i^\alpha)}$ has value $U^B_{Y(Q_i^{\alpha^{*,Mm}})} = U^{Mm}_{Y(Q_i^{\alpha^{*,Mm}})}$, since the minimax is obtained at the intersection M_i of $U_{Y(Q_i^\alpha,Yes)}$ and $U_{Y(Q_i^\alpha,No)}$.*

This means that we have : $\min_{X_i} U^B_{Y(Q_i^\alpha)} = U^B_{Y(Q_i^{\alpha^{*,B}})} \leq U^{Mm}_{Y(Q_i^{\alpha^{*,Mm}})} = \min_{X_i} U^{Mm}_{Y(Q_i^\alpha)}$, hence the expected uncertainty reduction with a Bayesian strategy is at least as high as the one obtained by the maximin strategy. However, in contrast with this latter, the Bayesian strategy does not offer guarantees about the uncertainty reduction, in the sense that the actual reduction may be lower than the expected one.

Also, while Proposition 1 means that $\alpha^{*,Mm}$ can be obtained by a dichotomic search, this cannot be done in general for the Bayesian strategy, which therefore requires heavier computations.

Example 3. *Consider the function $\Phi(x_1,x_2,x_3) = x_1x_2 - x_2x_3 + x_2$ with $X_1 = X_3 = [0.1,1]$ and $X_2 = [0,1]$. Figure 1 shows the various strategies for Q_2^α. We can see that the maximin, the Laplacian and Hurwicz's strategies recommend respectively $\alpha^{*,Mm} = 3/4$, $\alpha^{*,B} = 1/2$, and $\alpha^{*,H(\frac{1}{2})} = 0$. Another remark is that $U_{Y(Q_2^\alpha,Yes)}$ and $U_{Y(Q_2^\alpha,No)}$ are both linear. We will see in the next section that this is true for multi linear functions in general.*

3 QUERYING ON SPECIFIC FUNCTIONS

Here, we study what becomes of the previous strategies when applying them to specific functions. Indeed, finding an optimal strategy requires computing $U_{Y(Q_i^\alpha,Yes)}$, $U_{Y(Q_i^\alpha,No)}$ and their intersections, which comes down to finding bounds of Φ over various domains (see Eqs. (7)-(8)). It is therefore important to identify those sub-cases for which computations can be simplified. More precisely, we look at monotonic functions and multi linear functions, that are both of practical interest.

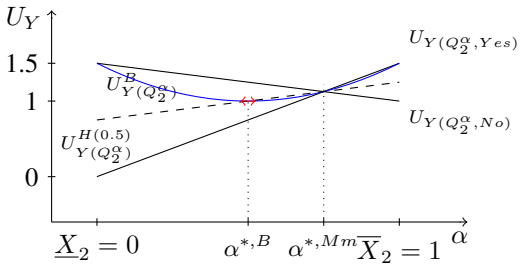

FIGURE 1 – Optimal recommendations of different query selection strategies.

3.1 MONOTONIC FUNCTIONS

Several application in diverse areas use monotonic functions, such as reliability analysis (Marichal, 2014), multi-criteria decision making (Grabisch and Labreuche, 2008), etc.

When considering such functions, either increasing or decreasing in each variable x_i, computations are greatly facilitated, as

$$U_Y = \Phi(\overline{X}_I, \underline{X}_{\overline{I}}) - \Phi(\underline{X}_I, \overline{X}_{\overline{I}}),$$

where I denotes the set of variables in which Φ is increasing, and $\overline{I}$ its complement.

Moreover, when Φ is locally monotonic [2] with respect to each argument i, its upper and lower bounds are reached on the vertices of the hypercube $\times_{i=1\ldots n} X_i$. Again, this may allow to reduce the computations involved in the calculation of U_Y.

3.2 MULTI LINEAR FUNCTIONS

A Multi-linear function over variables $x_1, \ldots, x_n$ is a polynomial form that can be written as

$$\Phi(x_1, ..., x_n) = \sum_{A \subseteq N} d_A \prod_{i \in A} x_i \qquad (9)$$

with $d_A \in \mathbb{R}$ some real-valued coefficients. Such functions play an important role in many AI applications. As any pseudo-Boolean function can be rewritten in this form (Hammer and Rudeanu, 1968), they concern all problems where pseudo-Boolean functions have a role, such as cooperative game theory (Owen, 1972), multi-criteria decision-making (Grabisch and Labreuche, 2003), combinatorial optimization (Yannakakis, 1991), reliability theory (Bhattacharjya and Deleris, 2012; Marichal, 2014), etc. Multi linear functions also play an important role in inferences of Bayesian networks or related models (Darwiche, 2003; de Campos and Cozman, 2004).

From a computational point of view, having Φ multi-linear presents different advantages. First, as ϕ is locally monotonic in each variable (fixing every variable values but one in Eq. (9) gives a linear function, which is either increasing or decreasing), we know that its upper and lower bounds are reached on vertices of $\times_{i=1\ldots n} X_i$. Second, provided $0 \notin X_i$, the maximin strategy will lead to a unique value $\alpha^{*,Mm}$, due to the fact that $U_{Y(Q_i^\alpha,Yes)}$, $U_{Y(Q_i^\alpha,No)}$ will be strictly increasing and decreasing, respectively (since bounds of Eq. (9) will be strictly monotone functions).

3.3 MULTI LINEAR MONOTONIC FUNCTIONS

Combining monotonicity and multi linear properties provide very interesting properties to compute our optimal strategies, and are still useful in several applications, such as reliability analysis that we use as a case study in the next section. The first property relates to the shape of $U_{Y(Q_i^\alpha,Yes)}$ and $U_{Y(Q_i^\alpha,No)}$

Proposition 2. *If Φ is a multi linear function monotonic in each variable, then for every $i \in N$, $U_{Y(Q_i^\alpha,Yes)}$ and $U_{Y(Q_i^\alpha,No)}$ are linear in α.*

Proof. *If Φ is monotonic in each variable, then in the first term of Eq.* (7)*, the maximum is reached on the upper bounds of each X_i, i.e., on $\overline{X}_j$ for all $j \in N_{-i}$ and $\overline{X}_i = \alpha$, while the lower bound is reached on $\underline{X}_j$ for all $j \in N$ (independent of α). The function Φ being linear in x_i, $\max \Phi(\mathbf{x})$ is therefore linear in α, and so is $U_{Y(Q_i^\alpha,Yes)}$. The same reasoning applies to Eq.* (8)*.*

This has several consequences on the computations of strategies :

- The maximin strategy recommends a unique query bound M_i in X_i, as $U_{Y(Q_i^\alpha,Yes)}$ and $U_{Y(Q_i^\alpha,No)}$ intersection will be a unique point ;
- Computing $U_{Y(Q_i^\alpha,Yes)}$ and $U_{Y(Q_i^\alpha,No)}$ will require only three computations, as they are linear (requiring each two evaluations) and as they have the same maximal value. Computing M_i then comes down to evaluate the intersection point of two lines ;
- Hurwicz's solution will be reached either at the endpoints of the interval X_i or will coincide with the maximin solution. The result follows from the fact that the convex combination of linear functions ($U_{Y(Q_i^\alpha,Yes)}$ and $U_{Y(Q_i^\alpha,No)}$) is also linear, and is therefore monotonic in α.

Furthermore, we have the following property regarding the Bayesian strategy :

Proposition 3. *If Φ is a multi linear function monotonic in each variable, the Bayesian strategy adopting a uniform distribution over X_i has a unique minimum α in the interior of X_i*

2. ϕ is locally monotone in x_i if, all other variables being fixed, it is either decreasing or increasing in x_i. Function ϕ of Example 1 is locally monotone in x_2, as it is either increasing or decreasing in x_2 once x_1 and x_3 are fixed.

Proof. *(sketch) Function $U^B_{Y(Q_i^\alpha)}$ is convex since it is the sum of the product of two linear functions of α, therefore it is quadratic and convex. In addition, it satisfies : $U^B_{Y(Q_i^{X_i})} = U^B_{Y(Q_i^{\overline{X}_i})}$.*

Since it can not be a constant function (the scenario $U_{Y(Q_i^\alpha, Yes)} = U_{Y(Q_i^\alpha, No)}$ for every $\alpha \in X_i$ does not occur for multi linear functions), its global minimum exists, is unique, and is reached inside the interval X_i.

4 APPLICATION IN RELIABILITY ANALYSIS

When systems are complex or newly designed, full system dependability data are often too expensive and/or difficult to obtain, making it impossible to directly estimate quantities of interest. The common approach to improve the estimation of such quantities is to focus on enhancing the state of knowledge at the component-level, where information is more likely to be available either via measurements or expert elicitation.

In this section, we illustrate how our elicitation model can be used to refine the state of knowledge at the component level in order to estimate the reliability of a system. We begin by recalling some basic elements related to systems reliability and briefly describe the mathematical properties of the system function. Then, we describe and discuss the results of the proposed elicitation procedure on simple yet common system architectures, to finish by a real-world example involving railway safety systems.

4.1 PRELIMINARIES ON RELIABILITY ESTIMATION

Consider a network S with n components indexed in $N = \{1, 2, ..., n\}$. We describe a static problem, i.e., we do not refer to time explicitly when describing the system behavior. Every component i is either operating or failing and its state is represented by a boolean variable e_i that associates 0 and 1 to the failed, working state, respectively. The system state is completely determined by the joint state of its components through the *structure function* Φ_S – a boolean function. For the majority of systems, forming the class of semi-coherent systems, the structure function satisfies these three conditions :
– Φ_s non-decreasing in each e_i
– $\Phi_s(0, 0, ..., 0) = 0$
– $\Phi_s(1, 1, ..., 1) = 1$.
In reality, the state of component can not be determined exactly, and the usual framework is to assume that is a random variable. The probability that the component is functioning is called the elementary reliability :

$$p_i = Pr(e_i = 1).$$

When the components are independent, i.e., when their state variables are stochastically independent, the reliability of the system :

$$R = Pr(\Phi_s(e_1, ...e_n) = 1),$$

can be determined from the reliability of its components via the *reliability function* Φ :

$$R = \Phi(p_1, ..., p_n), \tag{10}$$

which is the multi linear extension of Φ_s (Marichal, 2014) and so writes :

$$\Phi(p_1, ..., p_n) = \sum_{A \subseteq N} d_A \prod_{i \in A} p_i$$

where coefficients d_A are the Mobius transform of the mass function associated with the structure function [3]. In practice, the exact expression of the reliability function can be directly generated using the inclusion-exclusion formula (Lin et al., 1976) based on determining the minimal path set (i.e., the minimal set of components that must be in working state that guarantees the functioning of the system) and cut sets (the minimal set of components such that if all of them fail, the system is guaranteed to fail whatever the value of the others components).

Therefore, when facing a new system with ill-known probabilities, we have a multi linear function ϕ with interval-valued variables p_i, to which we can apply our previous findings.

4.2 CLASSICAL STRUCTURES

We first consider a bridge structure with 5 non redundant components. The reliability block diagram – a graphical depiction of the functional relationship between components – of this structure is given below :

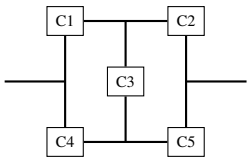

FIGURE 2 – Reliability block diagram of a series parallel system.

3. The Mobius transform of the mass function associated to Φ_s is given by :

$$d_A = \sum_{B \subseteq A} (-1)^{|A|-|B|} \Phi_s(B).$$

The system reliability is given by :

$$\begin{aligned} R \quad &= p_1p_2 + p_4p_5 + p_1p_3p_5 + p_2p_3p_4 \\ &-p_1p_2p_3p_4 - p_1p_2p_3p_5 - p_1p_2p_4p_5 - p_1p_3p_4p_5 \\ &-p_2p_3p_4p_5 + p_1p_2p_3p_4p_5. \end{aligned}$$

We assume the initial state of knowledge to be the following : $p_1 \in P_1 = [0.5, 0.92]$, $p_2 \in P_2 = [0.2, 0.9]$, $p_3 \in P_3 = [0.5, 0.9]$, $p_4 \in P_4 = [0.4, 0.8]$, and $p_5 \in P_5 = [0.4, 0.85]$. The system reliability ranges in the interval $[0.3, 0.97]$, so its initial uncertainty is 0.67.

We use the elicitation procedure to refine the state of knowledge over p_i ($i \in \{1, ..., 5\}$) via a sequence of queries on the elemental reliabilities. The objective is to reduce the system reliability uncertainty up to 0.05 (i.e., $s_0 = 0.05$), after which we stop asking questions.

To evaluate the efficiency of our procedure, we compare its performance to two basic strategies :

1. a random strategy that compares at each stage the reliability of component i, selected at random in N, with some random $\alpha \in P_i$. For the results to be significant, the performance of the strategy at each iteration is averaged over a high number (herein 1000) of runs.
2. a baseline strategy that asks at each stage about the most uncertain component, and the query bound is the midpoint of the largest interval (this strategy was referred to as the "halve largest Gap Strategy" in the context of preference elicitation (Boutilier et al., 2006)) :

$$Q^*_{Baseline} = \left(i^*, \frac{\overline{X}_{i^*} + \underline{X}_{i^*}}{2} \right)$$

where :

$$i^* = \arg\max_{i \in N} U_{X_i}.$$

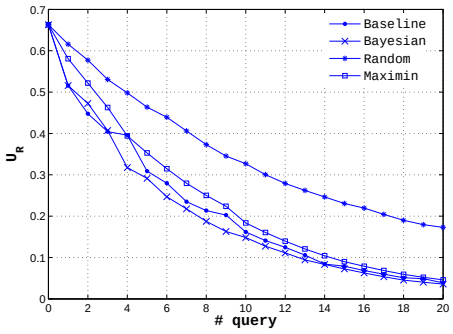

FIGURE 3 – U_R reduction using different selection strategies.

We implemented the elicitation procedure described in Section 2.3 assuming the true values to be the following : $p_1^* = 0.6$, $p_2^* = 0.7$, $p_3^* = 0.65$, $p_4^* = 0.7$, $p_5^* = 0.78$. Figure 3 shows the performance in terms of uncertainty reduction in R of our four strategies. The Bayesian slightly outperforms the baseline and the maximin strategies, but remains comparable to them, while all of them do much better than the random elicitation. In general, it takes twice the number of queries to the random strategy to reach the results of the other strategies (e.g., to divide uncertainty by half, it requires 10 questions for the random strategy, and about 5 for the others).

However, the performances of the Bayesian, maximin and baseline strategies highly depend on the initial situations. Figure 4 compares our previous experiment with another situation where the initial state of knowledge is very poor, i.e., a situation of near ignorance where $P_i = [0.1, 0.9]$ for all $i \in N$. Results for both scenarios differentiated by the color and the line style (blue continuous lines and black dashed lines for the first and the second scenarios, respectively). The most notable difference between the two scenarios concerns the maximin strategy. Indeed, its performance in the second scenario significantly departs from the non-random strategies (to which it was very close in the first scenario). The maximin strategy is in this case probably too cautious, missing potentially good opportunities to reduce the uncertainty.

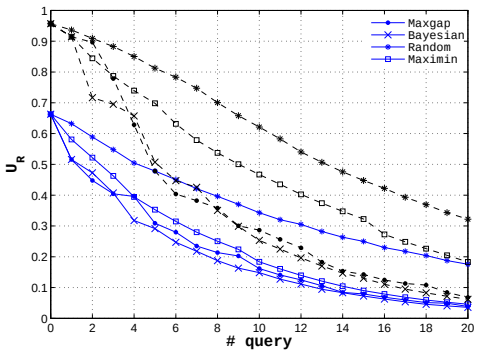

FIGURE 4 – Sensitivity of the performance of the selection strategies to the initial state of knowledge.

The good results of the baseline strategy for the bridge system are mainly due to the fact that every component is important in the system, hence gaining knowledge on any one of them reduces uncertainty in similar ways. This is not always true : consider a simple series parallel system composed of four independent and non-identical components (Fig. 5). The system reliability is :

$$R = p_1p_2p_3p_4 + p_1p_4 + p_2p_4 + p_3p_4 - p_1p_3p_4 - p_2p_3p_4 - p_1p_2p_4.$$

Let the initial state of knowledge on the elementary reliabilities be the following : $p_1 \in [0.01, 0.99]$, $p_2 \in [0.01, 0.99]$, $p_3 \in [0.97, 0.99]$, $p_4 \in [0.7, 0.9]$, and the true values be : $p_1 = 0.6$, $p_2 = 0.7$, $p_3 = 0.98$, $p_4 = 0.8$.

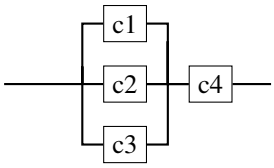

FIGURE 5 – Reliability block diagram of a series parallel system.

The results of the sequential elicitation procedure using the baseline strategy can be visualized in Figure 6 which plots the uncertainty on each component at every stage. A jump in the curve of component i at stage $k+1$ indicates that the k$^{\text{th}}$ optimal query inquired about that component, and its magnitude corresponds to the uncertainty reduction after the question has been answered. Note that up to the 6$^{\text{th}}$ question, the strategy inquired about the reliability of components 1 and 2, being the most uncertain.

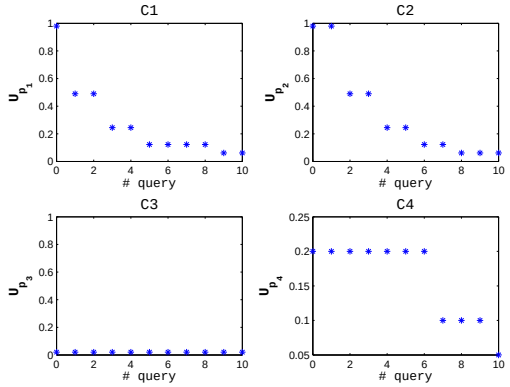

FIGURE 6 – Sequence of optimal components using the baseline strategy.

However, reducing uncertainty on components 1 and 2 does not reduce our global uncertainty, as shows Figure7. In this case the baseline strategy performs actually very bad, not only compared to the maximin strategy, but also to the random up to the 6$^{\text{th}}$ query. This is due to the fact that the baseline strategy does not consider the importance components have on the overall system reliability.

4.3 REAL CASE SYSTEM

Up to now, we considered simple structures with distinct (non-redundant) components. However, the majority of real systems are complex and redundant, i.e., some of their components are duplicated. Redundancy ensures a backup in case of failure of one of the critical parts, and aims at increasing the overall reliability of the system.

When a system has redundancies, its reliability function is no longer multi linear, and depending on the redundancy architecture (parallel, triple modular, etc.), it becomes polynomial in the reliability of the redundant components, while remaining linear in the others. We are concerned here with the study of this type of systems/functions.

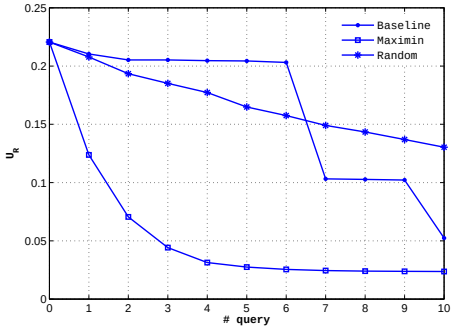

FIGURE 7 – Maximin, random, and baseline strategies for the case of series parallel system.

As a case study, we consider a real system used in the European railways traffic management system : the Radio Block Center system (RBC), whose role is to collect data about the position of trains and to provide movement authorisation (Flammini et al., 2006). Because of the relatively recent exploitation of the system, sufficient data to estimate the reliability of the RBC are lacking.

The RBC is composed of 5 different components, each of them being redundant. The architecture of the RBC is pictured in Figure 8, where the $2/3$ symbol means that the subsystem composed of components 5 works if and only if at least 2 out of the three components work.

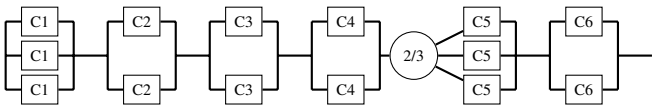

FIGURE 8 – Reliability block diagram of the RBC.

The reliability function can be written as :

$$R = (1-(1-p_1)^3)(1-(1-p_2)^2)(1-(1-p_3)^2)(1-(1-p_4)^2)p_{tmr};$$

with

$$p_{tmr} = (3p_5^2) - 2(p_5^3))(1-(1-p_6)^2).$$

We consider the case where some initial evidence suggests that the reliability of the RBC components ranges in $[0.5, 1]$, and that the true values are : $p_1 = 0.83$, $p_2 = 0.77$, $p_3 = 0.8$, $p_4 = 0.55$, $p_5 = 0.72$, $p_6 = 0.78$. Results of the query strategy are plotted in Figure 9. Here, the maximin strategy outperfoms the Bayesian one, which is consistent with Remark 1. The baseline does not do well and significant uncertainty reduction only occurs when asking about component 5, which is indeed the most important in this architecture.

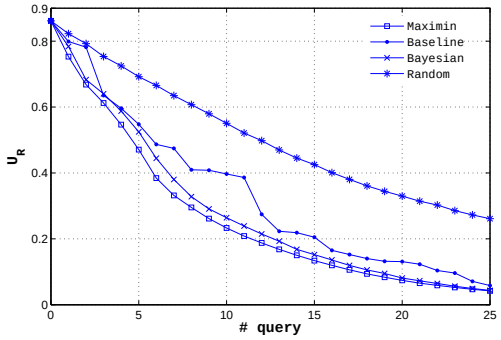

FIGURE 9 – Performance of the query strategies for the RBC system.

The computations involved in the elicitation procedure for this systems, and more complex systems in general, remain tractable as they only require optimization of polynomials and can still take advantage of the increasingness of function Φ. This makes our procedure of practical use in real-time elicitation involving real experts – this will be the object of a forthcoming work concerned with the estimation of the RBC reliability using expert elicitation.

5 CONCLUSION

In this paper, we addressed the problem of optimal expert elicitation when the goal is to reduce interval uncertainty. We described different optimal querying strategies to determine the best question to ask at each stage of the procedure, studied their computational costs, and illustrated their use in a common estimation problem in reliability analysis.

For the particular problem of interval uncertainty reduction using local bound queries, the optimal elicitation approach proves to be effective and computationally tractable, especially for the maximin approach. We also discussed some cases, such as monotonic and multi linear functions, for which these computations are even easier. In future works, we plan to consider (1) more general uncertainty models, such as belief functions (Shafer, 1976) or probability sets (Augustin et al., 2014) which are particularly appealing to model, e.g. non-completely reliable experts (allowing for instance to relax the assumption that the expert is an oracle) and (2) other types of queries formats and answers, such as comparative assessments.

The strategies we described in this paper are myopic. Such strategies offer natural advantages (any-time stop, computational easiness), yet may select a sequence of questions that are globally sub-optimal, despite being locally optimal. A natural extension of this work is then to address the sequential approach for selecting the optimal set of queries to ask, and compare it with the myopic method. Clearly, this includes dealing with a computationally challenging problem, given the multistage nature of the optimization task, as well as some potential difficulties when choosing the values of strategies (e.g., the over-cautious nature of maximin could lead to strategies with very low average performances).

Acknowledgments

This work was supported by the ANR (National Agency for Research) RECIF project (Reference : ANR-13-JS03-0007). It was carried out in the framework of the Labex MS2T.

References

Aspinall, W., Cooke, R., 2013. Quantifying scientific uncertainty from expert judgement elicitation. Risk and Uncertainty Assessment for Natural Hazards, 64.

Augustin, T., Coolen, F., de Cooman, G., Troffaes, M., 2014. Introduction to imprecise probabilities. John Wiley & Sons.

Benabbou, N., Perny, P., Viappiani, P., 2014. Incremental Elicitation of Choquet Capacities for Multicriteria Decision Making. In : Proceedings of ECAI'14. pp. 87–92.

Bhattacharjya, D., Deleris, L. A., 2012. From reliability block diagrams to fault tree circuits. Decision Analysis 9 (2), 128–137.

Boutilier, C., Brafman, R., Geib, C., Poole, D., 1997. A constraint-based approach to preference elicitation and decision making. In : AAAI Spring Symposium on Qualitative Decision Theory. Citeseer, pp. 19–28.

Boutilier, C., Filmus, Y., Oren, J., 2013. Efficient vote elicitation under candidate uncertainty. IJCAI.

Boutilier, C., Patrascu, R., Poupart, P., Schuurmans, D., 2006. Constraint-based optimization and utility elicitation using the minimax decision criterion. Artificial Intelligence 170 (8-9), 686–713.

Braziunas, D., Boutilier, C., 2007. Minimax regret based elicitation of generalized additive utilities. In : UAI. pp. 25–32.

Chajewska, U., Koller, D.and Parr, R., 2000. Making rational decisions using adaptive utility elicitation. In : AAAI/IAAI. pp. 363–369.

Cooke, R., 1991. Experts Uncertainty. Oxford University Press.

Curtis, A., Wood, R., 2004. Optimal elicitation of probabilistic information from experts. Geological Society, London, Special Publications 239 (1), 127–145.

Darwiche, A., 2003. A differential approach to inference in bayesian networks. Journal of the ACM (JACM) 50 (3), 280–305.

de Campos, C., Cozman, F., 2004. Inference in credal networks using multilinear programming. In : Proceedings of the Second Starting AI Researcher Symposium. pp. 50–61.

Flammini, F., Marrone, S., Mazzocca, N., Vittorini, V., 2006. Modelling system reliability aspects of ertms/etcs by fault trees and bayesian networks. In : Proc. European Safety and Reliability Conference, ESREL. pp. 2675–2683.

Grabisch, M., Labreuche, C., 2003. On the extension of pseudo-boolean functions for the aggregation of interacting criteria. European Journal of Operational Research.

Grabisch, M., Labreuche, C., 2008. A decade of application of the choquet and sugeno integrals in multi-criteria decision aid. 4OR 6 (1), 1–44.

Guerin, J., Allen, T. E., Goldsmith, J., 2013. Learning cp-net preferences online from user queries. In : Algorithmic Decision Theory. Springer, pp. 208–220.

Hammer, P., Rudeanu, S., 1968. Boolean Methods in Operations Research and Related Areas. Springer, Berlin Heidelberg, New York.

Keeney, R. L., Raiffa, H., Rajala, D., 1979. Decisions with multiple objectives : Preferences and value trade-offs. Systems, Man and Cybernetics, IEEE Transactions on 9 (7), 403–403.

Lin, P., Leon, B., Huang, T., 1976. A new algorithm for symbolic system reliability analysis. Reliability, IEEE Transactions on R-25 (1), 2–15.

Marichal, J.-L., 2014. Structure functions and minimal path sets. arXiv preprint arXiv :1401.0803.

Owen, G., 1972. Multilinear extensions of games. Management Sciences 18, 64–79.

Shafer, G., 1976. A Mathematical Theory of Evidence. Princeton University Press, Princeton, N. J.

Spetzler, C., Stael von Holstein, C., 1975. Probability encoding in decision analysis. Management Science, 340–358.

Viappiani, P., Kroer, C., 2013. Robuts optimization of recommendation sets with the maximin utility criterion. Proceedings of the Third Algorithmic Decision Theory International Conference, 411–424.

Wang, T., Boutilier, C., 2003. Incremental utility elicitation with the minimax regret decision criterion. In : IJCAI. pp. 309–318.

Winkler, R., 1969. Scoring rules and the evaluation of probability assessors. Journal of the American Statistical Association 64 (327), 1073–1078.

Yannakakis, M., 1991. Expressing combinatorial optimization problems by linear programs. Journal of Computer and System Sciences 43 (3), 441 – 466.