

---

# A Web-based Tool for Expert Elicitation in Distributed Teams

---

**Carlo Spaccasassi**  
IBM Dublin Research Lab  
Damastown, Dublin, D15  
Ireland  
spaccasa@ie.ibm.com

**Lea Deleris**  
IBM Dublin Research Lab  
Damastown, Dublin, D15  
Ireland  
lea.deleris@ie.ibm.com

## Abstract

We present in this paper a web-based tool developed to enable expert elicitation of the probabilities associated with a Bayesian Network. The motivation behind this tool is to enable assessment of probabilities from a distributed team of experts when face-to-face elicitation is not an option, for instance because of time and budget constraints. In addition to the ability to customize surveys, the tool provides support for both quantitative and qualitative elicitation, and offers administrative features such as elicitation surveys management and probability aggregation.

## 1 Introduction

There is a thriving research community that studies techniques for learning the structure and parameters of a belief network from data [1]. However, when there is no relevant data available, or any literature to guide the construction of the model, the network must be elicited from the individuals whose beliefs are being captured - such a person is often referred to as the *domain expert*, or simply expert. Both the structure and the parameters of a belief network need to be elicited. Often it is easier to construct the *structure* of a belief network, as compared to eliciting the *parameters*, i.e. the conditional probabilities [2, 3, 4, 5]. We focus in this paper on the subject of parameter elicitation, assuming that the structure of the network has already been ascertained.

Best practice in terms of parameter elicitation is based on face-to-face interviews of the expert by a trained analyst (or knowledge engineer) [6, 7, 8]. However, situations arise where such an approach is not feasible, mostly because of time and budget constraints. This is especially salient in projects with a distributed team of

experts, which as Bayesian modeling gains popularity are more likely to arise than in the past [9].

Consider the following real-world example. We undertook a project focused on understanding variability in the performance of a specific human resource process and elected to use a Bayesian network as our modeling framework. The domain experts were regular employees acting as experts, they were scattered across the world and spanned different domains of expertise. We did not have the possibility of undertaking face-to-face sessions and opted for replacing them with phone interviews. The structural definition of the model, identifying the variables and inter-dependence, did not yield many difficulties nor complaints from the experts. By contrast, the quantitative phase proved time consuming and generated significant frustration on both sides (analysts and experts). In particular, our efforts were hampered by (i) the time difference leading to early or late at night sessions for either the expert or the analyst and (ii) the time pressure on the experts because of the analyst waiting on the phone for them to provide an answer. The main challenge however was to have experts understand the format of the conditional probability table (CPTs). Overall, we concluded that phone elicitation was not an adequate support for remote parameter elicitation and that eliciting probabilities directly in the CPT created unnecessary cognitive burden.

The risk elicitation tool that we present here aims at addressing those concerns. We opted for a web-based tool, whose asynchronous feature enables more comfortable time management of the elicitation process from experts side (albeit less control for the analyst). An advantage of the web-based set up is the ability for the analyst to centrally manage the elicitation surveys. While we recognize that web-based approaches are second-best to face-to-face elicitation, we feel that such a tool would enable wider adoption of Bayesian models in settings where face-to-face elicitation is unlikely.

The remainder of the paper is organized as follows. In Section 2, we review the literature related to probability elicitation in Bayesian networks. Section 3 provides an in-depth description of the Risk Elicitation tool. Finally Section 4 discusses related research endeavors.

## 2 Expert Elicitation in Belief Networks

The process of eliciting probabilities from experts is known to be affected by numerous cognitive biases, such as overconfidence and anchoring effects [10]. When eliciting probabilities in the context of a belief network, additional practical challenges must be considered [11].

One particular problem lies with the number of parameters that have to be elicited from the experts, which leads to long and tiring elicitation sessions and sometimes inconsistent and approximate answers. To alleviate such problems, the analyst often resorts to making assumptions about the conditional relationships that reduce the number of parameters to be elicited by parameterizing the network structures using NOISY-OR and NOISY-MAX models (see for instance [2, 12]). This is in fact an option that we will provide in the next version of our tool.

As we mentioned in the introduction, another challenge associated with elicitation in Bayesian networks is the problem for the expert to understand the structure of a conditional probability table. While considering scenarios is fairly intuitive, understanding which entry corresponds to which scenario can be unnecessarily confusing. Efforts have thus been made to improve the probability entry interface in probability elicitation tools [13, 14]. Our tool integrates findings from this stream of research, by asking simple text questions corresponding to each cell of the CPT and by grouping all assessments corresponding to the same scenario together (although our support does not enable us to show them all at once but simply sequentially). Indeed, previous research has shown that presenting all conditioning cases for a node together during elicitation reduces the effect of biases [5].

Finally, the need to provide precise numerical answers is considered an additional cognitive obstacle for experts. One solution to address this problem is to present the elicitation scale with verbal and numerical anchors [15, 5, 16]. We included such findings into the design of our tool, enabling analyst to ask questions in a qualitative manner. Another solution is to elicit qualitative knowledge from experts, for instance by asking them to provide a partial order of the probabilities and leveraging limited data whenever available

[17].

## 3 Description of the Tool

### 3.1 Overview

The Risk Elicitation tool is a web-based application that offers both (i) an interactive web interface through which parameter elicitation surveys themselves can be answered and automatically collected, and (ii) support for survey management. The tool can be freely accessed from the Internet; any web browser with Adobe's Flash Player 10 [18] installed will be able to run it. Given its web availability, the Risk Elicitation tool is virtually always available. Moreover, interviewees can complete a survey with little external help, pause and resume the survey at a later time, thus further relaxing the need to coordinate interviewers and interviewees.

We distinguish two classes of users of the Risk Elicitation tool: *analysts* and *domain experts*. In the following sections we describe the main use cases of the tool setting up elicitation surveys (Analyst), answering a survey (Expert) and collecting and aggregating results (Analyst). We also provide at the end of this section a description of the architectural set up along with a short discussion of the technical challenges that we met.

### 3.2 Setting up Elicitation Surveys

As mentioned earlier, we assume that the starting point of the process is a Bayesian network whose structure is fully defined, including clear description of nodes and associated states. The first step for the analyst is therefore to load his Bayesian Network file on the Risk Elicitation tool. The tool will automatically generate a sample survey, which the analyst can further customize. The second step for the analyst is to create a user account for each expert. Experts, having various domains of expertise, may not be qualified to provide information for all the nodes in the Bayesian network. To address that situation, the analyst can define roles and associate a subset of the nodes to each role. Each expert can then be associated to one or several roles and will only be asked questions on the Bayesian nodes pertaining his/her role(s)<sup>1</sup>.

The main features of the tool that enable survey setup are:

**BBN Import** The Risk Elicitation tool enables analysts to create a personalized survey of the BBN

---

<sup>1</sup>In the remainder of this paper we will refer to expert and analyst as *he*.

**Ctx:** Consider the variable "Infrastructure Availability"

**Severe Earthquake is Occur**

**Pandemic Occurrence is Occur**

**More details**

**Infrastructure Availability's local network**

**Severe Earthquake    Pandemic Occurrence**

**Infrastructure Availability**

---

**Q:** Please use the sliders on the wheel below to provide the probabilities of the different states of Infrastructure Availability (High, Medium, Low)

State	Probability
High	39.53%
Medium	42.67%
Low	17.80%

This implies that High is 2.22 times more likely to occur than Low.

This implies that Medium is 2.40 times more likely to occur than Low.

**Please select a confidence level:**

High: 5%   
 Medium: 10%   
 Low: 25%

**✕ Please provide your comments below:**

Question 1 of 4    **Back**    **Next**

Figure 1: Expert Elicitation page for a quantitative question.

they want to elicit. The BBN can be submitted from the Risk Elicitation tool to a server that automatically generates a survey template. The template can then be customized by the analyst, who can perform the following modifications:

- Provide descriptive details on the Bayesian networks, its nodes and its states; add analyst notes to specific questions,
- Choose how to elicit node, whether quantitatively or qualitatively,
- Define, for the qualitative questions, the possible answers and relative numerical ranges (which we call *calibrations*),
- Customize the question texts,
- Choose whether to ask experts about their confidence level,
- Assign an order to the elicitation process (to control in which order nodes are elicited).

At the moment we only support the GeNIe file format, but our tool can be easily extended to other formats. We have developed our own format for

Bayesian networks, to which the GeNIe file format is translated during the template generation.

**User Management** In the Administration section, analysts can register experts to the Risk Elicitation tool and assign them roles. The analyst is presented with a classic user management console, where he can add, delete and update both user accounts and the roles they play in an elicitation survey. Whenever a user account is created, the tool generates an automatic email, that the analyst can further customize and send to the expert, presenting him his credentials to access the tool and the survey he has been assigned.

### 3.3 Expert Elicitation

After an expert has been notified of his account credentials, he can access the Risk Elicitation tool. Upon logging in, he can select one of the surveys and roles he has been assigned to. At that point, he is offered the option of reviewing a short tutorial of the tool. Moving to the survey answering, he is presented with questions

**Pandemic Occurrence Occurrence Page**


**Ctx:** Imagine the following situation

---

**Comments from the survey creator**

**More details**

**Pandemic Occurrence Node's local network**



Pandemic Occurrence Node

**Q:** Please select a range for the probabilities of the different states of Pandemic Occurrence Node (Occur, Does Not Occur)

Please provide an assessment for the probability that **Pandemic Occurrence Node** takes the value **Occur**:

<input type="radio"/> Low: 33% chance or below	<b>Confidence level:</b>
<input type="radio"/> Medium: from 33% to 66% chance	<input checked="" type="radio"/> High: 5%
<input type="radio"/> High: 66% chance or above	<input type="radio"/> Medium: 10%
	<input type="radio"/> Low: 25%

Please provide an assessment for the probability that **Pandemic Occurrence Node** takes the value **Does Not Occur**:

<input type="radio"/> Low: 33% chance or below	<b>Confidence level:</b>
<input type="radio"/> Medium: from 33% to 66% chance	<input checked="" type="radio"/> High: 5%
<input type="radio"/> High: 66% chance or above	<input type="radio"/> Medium: 10%
	<input type="radio"/> Low: 25%

[Revisit Calibration for this question](#)

**X** Please provide your comments below:

Question 1 of 1 [Back](#) [Next](#)

Figure 2: An example of qualitative question.

for each relevant node of the Bayesian network. They can ask for either quantitative or qualitative answers. When the survey is complete, the expert can submit the survey on the Risk Elicitation tool and exit.

Expert elicitation is supported by the following features:

### Quantitative and Qualitative Elicitation

Probabilities can be elicited through either *quantitative* or *qualitative* questions. Quantitative questions ask experts to state exact probabilities, using a pie chart for discrete nodes. As shown in Figure 1, each slice of the pie represents a state of the Bayesian node with its associated probability. Users can drag the pie chart edges to provide their estimates of the node being currently evaluated, given that the scenarios defined in the context pane (parent nodes and states), at the top left corner of the question page. We also provide direct feedback about the implied odd ratios on the right side of the pie chart, as some situations may be more suited to thinking about relative chances. The map in the top right corner shows the local network topology for the node being elicited. The full Bayesian Network is also available in the *Road Map* tab on the left-hand side.

*Qualitative* questions do not elicit exact probabil-

ities but ranges of probabilities. As shown in Figure 2, experts are offered a set of labeled ranges, called *calibrations*, and can select the calibration that best describes the probability of a node being in a state, given the conditions expressed in the Context pane. Calibrations are initially defined by analysts at BBN Import time, both in terms of labels and numerical range. However, experts have the ability to modify the numerical values of ranges from the tool itself if they feel they are not appropriate for the specific question.

**Summary Tables** There are as many questions for each node as parent state configurations. After all questions for a node have been answered, the expert is shown a summary table that provides a report of all the answers they have given (see Figure 3). This is in fact the conditional probability built from the answers provided. However, at this point the expert has been actively involved in building it from the ground up and should not be as confused by the structure as if we had presented it upfront. The summary table enables to compare answers across scenarios. If the expert wants to change any of the input, he can navigate back to the associated question by clicking on the related cell in the summary table, as shown in Figure 3. When the expert is satisfied with his answers, he can save and proceed to either answer questions about another node, or submit the survey if all nodes have been answered.

**Confidence** For each question/node, the expert can provide an indication of his confidence in his answer (provided the analyst has enabled this feature). At this point, confidence indication is qualitative (Low/Medium/High) but could be further defined in terms of notional sample space for instance. Confidence information can be used during aggregation, to modify the weight an answer has, or to provide a threshold to filter out answers (e.g. consider only high confidence answers).

**Comment** For each question, the expert has the opportunity to provide a comment through an apposite collapsible text area, placed below the question itself. One use of the comment section is to provide details about understanding of a node description or state or to specify an implicit assumption that the expert has made when providing answers.

### 3.4 Gathering information

After setting up surveys and notifying experts, the analyst can use the Status section of the tool to check

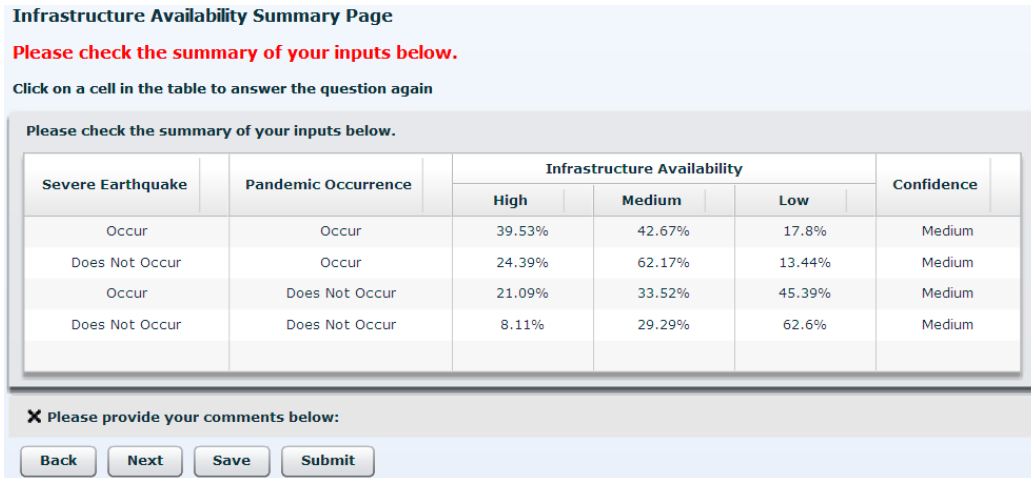


Figure 3: Summary page for the questions elicited in Figure 1.

on the progress of the elicitation process. He is provided with a summary of how many surveys have been completed. From the same section, experts can be reminded to complete their survey by an automatically generated email. Once enough surveys have been completed, the analyst has the option to aggregate expert answers and export a file of the Bayesian network populated with the aggregated values.

The main mechanisms to enable gathering and aggregation of answers are:

**Surveys Monitoring** Analysts can monitor the advancement of survey completion from a dedicated section, called Status Tab. The Status Tab reports which experts have completed their surveys and when, which surveys have not been submitted yet and which experts have been reminded to finish the survey. To remind an expert to complete his survey, an automatic mailing system is provided to automatically generate and send email reminders to the interested parties. Generated email kindly remind experts of which surveys they have been assigned, the role they play into it, their account details in case they forgot and a link to the tool. Analysts can also customize the generated email before sending it from the tool itself, as shown in Figure 4.

**Probability Aggregation** After all surveys have been completed, an analyst may need to aggregate the answers provided by the experts. We currently support two methods of aggregation: *linear opinion pool* and *logarithmic opinion pool* [19]. The analyst can control the aggregation method by assigning a weight to each expert, to credit some experts more importance. The tool goes through all completed surveys, collects the prob-

abilities elicited by experts and aggregates them using the method and weights specified by the analyst. Given that qualitative questions do not provide an exact number but a range, we take the midpoint of each range as the representative of the range (while acknowledging that this is a rather simple approach which we will refine in later versions of the tool).

Aggregated values are used to populate the original Bayesian Network file imported in the tool. The analyst can then export the aggregated BBN on his computer.

### 3.5 Implementation Details

The tool employs a classic two-tiers architecture, with a web application developed on top of IBM's Websphere Application Server 6.1 [20] and a Flash client built with Adobe's Flex Builder 3 [18]. We have employed a Model-Driven Architecture approach [21] to develop the tool, following the standard Model-View-Control pattern, where the view is the Flash client, most of the controls are in the web server and the model is the survey itself, exchanged and modify by both server and clients. Communication is handled by web services using JAX-RPC [22].

Surveys data has been modeled using the Eclipse Modeling Framework (EMF) [23]. We first designed an abstract, graphical representation of the data that the survey needed to capture in EMF. The resulting representation, or model, is similar to a UML Class diagram. Code to manipulate and also persist the model is automatically generated from the model and taken care of by EMF.

EMF does not support natively Actionscript, Adobe's programming language: EMF's standard tools cannot

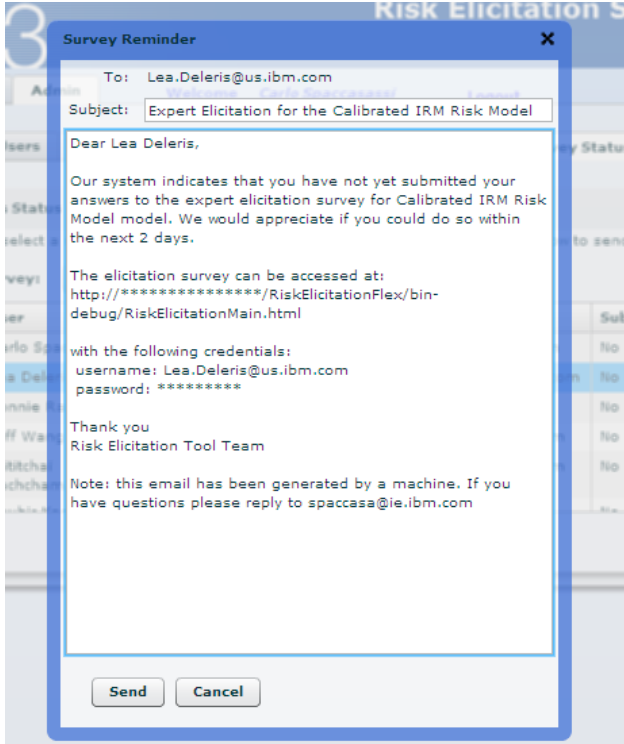


Figure 4: Customizable email templates from the administration console.

generate model manipulation code automatically for it. To address this problem, we bridged EMF to a Web Service definition file (WSDL). We first exported EMF’s models to an XML Schema, which we imported into the WSDL file. Adobe’s Flex can then generate code from the WSDL file both to communicate with the server and to access the model.

Communication points between server and client are also generated from the WSDL file. Extensions and modifications to either the model or the communication points, on the server and client side, were handled automatically by either EMF or Flex, handling manual error-prone tasks and saving development time.

Finally, we import and export BBN files written in the SMILE/GeNIe format [24]. EMF automatically manages GeNIe file loading and saving, using the XML Schema definition which is publicly available. The GeNIe files are then converted to an internal EMF model designed to ease BBN manipulation.

## 4 Related Research

In this section, we briefly discuss some of the research questions that have arisen from the development of the risk elicitation tool. In particular, we have focused on the effect on the elicitation process of the order in

which the nodes are presented. Because of the web-based feature of our tool, we have more freedom in determining the order than traditional face-to-face approaches.

### 4.1 Experiencing Different Orders

We considered the following question: *Does the parameter elicitation ordering in belief networks even matter to a domain expert?* To answer this question, we explored the relationship between node ordering and user-friendliness of the elicitation process in an experimental setting. Specifically, three different node orderings for the same belief network were considered: two ‘top-down’ and one ‘bottom-up’ ordering, with parameter elicitation performed using the risk elicitation tool described in this paper. Around seventy Stanford University graduate students were asked to elicit a belief network with six nodes on the subject of getting a job immediately after their studies; they were split into approximately three equal groups, one group for each order. The top-down orders presented questions to elicit parameters of parent nodes before children nodes, while the bottom-up order visited children before parents.

In this particular experiment, there was no drop-out - all subjects completed the elicitation process, perhaps due to the small size of the network and the incentive of extra class credit (which was only granted for complete assessments). Along with the Web-based elicitation survey, the students also responded to a short survey requesting feedback about the elicitation process and the corresponding tool. The results did indicate that the order in which the nodes are presented affects not only how comfortable experts claim to be with the process, but also the time required to elicit the parameters. In particular, there was a significant difference between the orders with regard to user-friendliness based on the survey responses. For the two top-down orders, hardly any of the subjects felt that the order was confusing, compared to 23% for the bottom-up order. Moreover, the average time to complete the elicitation was lower for the two top-down orders as compared to the bottom-up order. The two top-down orders differed in survey completion time: an average of 400 seconds with a standard deviation of 170 seconds for the first one, against an average of 500 seconds with a standard deviation of 400 seconds for the second one.

### 4.2 Ordering Mathematically

In a separate study, we explored the problem of determining, for a particular belief network whose structure is known, the optimal order in which the parameters

of the network should be elicited. Our objective in determining the order is to maximize information acquisition. While the order of the elicitation process is irrelevant if all nodes are elicited and if experts are able to provide their true beliefs, we believe that new trends in belief network modeling make these assumptions questionable. When only a subset of the nodes may be elicited or when answers can be noisy, it is necessary to devise an ordering strategy that seeks to salvage as much information as possible.

We therefore developed an analytical method for determining the optimal order for eliciting these probabilities, where optimality is defined as shortest distance to the true distribution (on which we have a prior). We considered the case where experts may drop out of the elicitation process and modeled the problem through a myopic approach.

For the case of uniform Dirichlet priors, we show that the ‘bottom up’ elicitation heuristic can be optimal. For other priors, we showed that the optimal order often depends on which variables are of primary interest to the analyst (whether all the nodes in the network or a subset, as is often the case in risk analytic applications).

The orderings resulting from the methods proposed in that model are driven solely by analytical concerns, and do not consider the user-friendliness of the elicitation process. In practice, as we discussed in the previous section, different orderings can impact the perceived difficulty of the process, thereby making the elicitation of complete and accurate beliefs more difficult. These results further motivate the need to investigate the consequence of forcing a possibly unnatural ordering upon experts and to assess whether the ‘information gain’ from an analytical perspective is worth the ‘cost’ in practice, i.e. in terms of the amount of confusion, fatigue and increased imprecision. More generally, empirical research to investigate how experts actually react to different orders is an important topic, similar to the empirical work on understanding how experts actually feel about different probability elicitation tools [15, 14]. The tool presented in this paper could be a useful support for such endeavors.

### 4.3 Comparison with existing web-based tools

Finally, we compare our tool to two existing web-based tools for risk elicitation pointed out by reviewers: BayesiaLab [25] and the Elicitation Tool from ACERA, the Australian Center of Excellence of Risk Analysis, described in [26].

BayesiaLab, a commercial product developed by Bayesia, provides an integrated environment for work-

ing with Bayesian networks. It supports many features, such as BBN modeling, BBN learning from data, and elicitation. With respect to elicitation, analysts can create a profile for each expert, select the portions of a variables’ CPT to be elicited, and send this information to a web server over the Internet. The web server generates surveys to elicit probabilities quantitatively, using a slider bar to capture expert input. Experts can also provide a level of confidence in an answer, expressed as a percentage, along with additional comments. In comparison with our tool, many features are similar: both tools provide expert profile management, on-line surveys, and survey import and export. Our tool, however, allows for both quantitative and qualitative elicitation of probabilities. The elicitation formats are different as well: our tool uses pie charts to capture quantitative probabilistic information for discrete random variables and a slider bar (expressed as a percentage difference from baseline) to capture impact of a factor on a (continuous-valued) metric under a specified scenario. Additionally, our tool allows analysts to fully customize surveys and aggregate by one of several algorithms. We also provide experts with additional contextual information, including a local and global map of the Bayesian network, a tutorial, the description of each node and state in the Bayesian network, along with analysts’ comments.

The Elicitation Tool from ACERA is quite different from both our tool and BayesiaLab, in that it is an online questionnaire to directly elicit estimates of risks. Questions are open-ended. An example is: “Will DAGGRE win?”. When answering a question, users need to provide four numerical estimates in an HTML form: the lowest estimate, the highest estimate, the best estimate and a confidence level. A graphical representation of the estimates is displayed and the user can submit the answers. After submission, the tool displays a selection of answers from users who have already completed the survey. The user is given the chance to review his own answers in light of this new input and submit again. In contrast, our tool allows review of only the expert’s own answers, as shown in Summary Pages, and are tailored for Bayesian networks where the goal is to elicit (conditional) probabilistic information. To help experts frame the context of the question, we provide additional information, such as the Bayesian network’s local map and network description. ACERA’s tool does not seem to be tied to Bayesian networks, so less contextual information is required in this case.

## 5 Conclusion

In this paper, we describe a web-based expert elicitation tool for Bayesian network models that is especially

relevant for the management of distributed teams of experts. We focus especially on facilitating the understanding of a conditional probability table by asking each entry separately and in a textual format. The tool enables the management of the survey administration cycle, from the customization of the survey and the creation of roles (associated with a subset of the network) to the monitoring of progress from experts and aggregation of results.

While we have implemented several best-practices from the elicitation literature, we also have identified various directions for further development. One simple extension will consist in allowing for NOISY-OR and NOISY-MAX parameterization. Going further, we would like to more strongly encourage for qualitative elicitation, asking for orders of magnitudes for instance, or if limited data was available following the relative order procedure suggested by [17]. In fact, for cases where partial data is available, one could also consider providing feedback to the expert directly during the elicitation session [27]. Finally, we have started providing support for utility/value nodes but so far in a coarse manner. Initially, experts are asked to identify a parent states configuration for which they are comfortable with providing an exact estimate of utility. We call this configuration *base case*. For non-base case configurations, experts only need to specify how much in percentage the utility of the node differs from the base case.

### Acknowledgements

This work is partially supported by fundings from IDA Ireland (Industrial Development Agency).

### References

- [1] D. Heckerman. A tutorial on learning with bayesian networks. In M. Jordan, editor, *Learning in Graphical Models*. Kluwer, Netherlands, 1998.
- [2] M. Henrion. Some practical issues in constructing belief networks. In L. Kanal M. Henrion, R. Shachter and J. Lemmer, editors, *Proceedings of the Fifth Conference on Uncertainty in Artificial Intelligence*, pages 161–173. Elsevier Science, New York, NY, 1989.
- [3] M. Druzdzel and L. van der Gaag. Elicitation of probabilities for belief networks: Combining qualitative and quantitative information. In P. Besnard and S. Hanks, editors, *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 141–148. Morgan Kaufman, San Francisco, CA, 1995.
- [4] M. Druzdzel and L. van der Gaag. Building probabilistic networks: Where do the numbers come from? *IEEE Transactions on Knowledge and Data Engineering*, 12(4):481–486, 2000.
- [5] L. van der Gaag, S. Renooij, C. Witteman, B. Aleman, and B. Taal. How to elicit many probabilities. In K. Laskey and H. Prade, editors, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 647–665. Morgan Kaufman, San Francisco, CA, 1999.
- [6] C. Spetzler and C. von Holstein. Probability encoding in decision analysis. *Management Science*, 22(3):340–358, 1975.
- [7] M. Merkhofer. Quantifying judgmental uncertainty: Methodology, experiences and insights. *IEEE Transactions on Systems, Man and Cybernetics*, 17(5):741–752, 1987.
- [8] R.L. Keeney and D. von Winterfeldt. Eliciting probabilities from experts in complex technical problems. *Engineering Management, IEEE Transactions on*, 38(3):191–201, aug 1991.
- [9] Sandra Hoffmann, Paul Fishbeck, Alan Krupnick, and Michael McWilliams. Elicitation from large, heterogeneous expert panels: Using multiple uncertainty measures to characterize information quality for decision analysis. *Decision Analysis*, 4(2):91–109, 2007.
- [10] D. Kahneman, P. Slovic, and A. Tversky. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press, 1982.
- [11] S. Renooij. Probability elicitation for belief networks: Issues to consider. *The Knowledge Engineering Review*, 16(3):255–269, 2001.
- [12] Adam Zagorecki and Marek Druzdzel. Knowledge engineering for bayesian networks: How common are noisy-max distributions in practice? In *Proceeding of the 2006 conference on ECAI 2006: 17th European Conference on Artificial Intelligence August 29 – September 1, 2006, Riva del Garda, Italy*, pages 482–486, Amsterdam, The Netherlands, The Netherlands, 2006. IOS Press.
- [13] Hope Nicholson Korb, L. R. Hope, A. E. Nicholson, and K. B. Korb. Knowledge engineering tools for probability elicitation. Technical report, 2002.
- [14] H. Wang and M. Druzdzel. User interface tools for navigation in conditional probability tables and elicitation of probabilities in bayesian networks. In C. Boutilier and M. Goldszmidt, editors, *Proceedings of the Sixteenth Conference on*



- Uncertainty in Artificial Intelligence*, pages 617–625. Morgan Kaufmann, San Francisco, CA, 2000.
- [15] S. Renooij and C. Witteman. Talking probabilities: Communicating probabilistic information with words and numbers. *International Journal of Approximate Reasoning*, 22:169–194, 1999.
- [16] F. Fooladvandi, C. Brax, P. Gustavsson, and M. Fredin. Signature-based activity detection based on bayesian networks acquired from expert knowledge. In *Information Fusion, 2009. FUSION '09. 12th International Conference on*, pages 436–443, july 2009.
- [17] E. M. Helsper, L. C. van der Gaag, A. J. Feelders, W. L. A. Loeffen, P. L. Geenen, and A. R. W. Elbers. Bringing order into bayesian-network construction. In *Proceedings of the 3rd international conference on Knowledge capture, K-CAP '05*, pages 121–128, New York, NY, USA, 2005. ACM.
- [18] Jeff Tapper, Michael Labriola, Matthew Boles, and James Talbot. *Adobe Flex 3: training from the source*. Adobe Press, first edition, 2008.
- [19] Robert T. Clemen and Robert L. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19:187–203, 1999.
- [20] E.N. Herness, R.H. High, and J.R. McGee. Web-sphere application server: a foundation for on demand computing. *IBM Syst. J.*, 43(2):213–237, April 2004.
- [21] Richard Soley and the OMG Staff Strategy Group. Model-driven architecture. <http://www.omg.org/~soley/mda.html>, 2000.
- [22] Roberto Chinnici. Java APIs for XML based RPC (JSR 101). <http://jcp.org/aboutJava/communityprocess/first/jsr101/>, October 28, 2003.
- [23] Frank Budinsky, Stephen A. Brodsky, and Ed Merks. *Eclipse Modeling Framework*. Pearson Education, 2003.
- [24] Marek J. Druzdel. Smile: Structural modeling, inference, and learning engine and genie: a development environment for graphical decision-theoretic models. In *Proceedings of the sixteenth national conference on Artificial intelligence and the eleventh Innovative applications of artificial intelligence conference innovative applications of artificial intelligence, AAI '99/IAAI '99*, pages 902–903, Menlo Park, CA, USA, 1999. American Association for Artificial Intelligence.
- [25] Bayesia. Bayesialab. <http://www.bayesia.com/en/products/bayesialab.php>.
- [26] Andrew Speirs-Bridge, Fiona Fidler, Marissa McBride, Louisa Flander, Geoff Cumming, and Mark Burgman. Reducing overconfidence in the interval judgments of experts. *Risk Analysis*, 30(3):512–23, 2010.
- [27] A. H. Lau and T. Y. Leong. Probes: a framework for probability elicitation from experts. In *Proceedings of the AMIA Symposium*. American Medical Informatics Association, 1999.