

Morris, David E. and Oakley, Jeremy E. and Crowe, John A. (2014) A web-based tool for eliciting probability distributions from experts. *Environmental Modelling and Software*, 52 . pp. 1-4. ISSN 1873-6726

Access from the University of Nottingham repository:

<http://eprints.nottingham.ac.uk/35002/1/1-s2.0-S1364815213002533-main.pdf>

Copyright and reuse:


The Nottingham ePrints service makes this work by researchers of the University of Nottingham available open access under the following conditions.

This article is made available under the Creative Commons Attribution licence and may be reused according to the conditions of the licence. For more details see:
<http://creativecommons.org/licenses/by/2.5/>

A note on versions:

The version presented here may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the repository url above for details on accessing the published version and note that access may require a subscription.

For more information, please contact eprints@nottingham.ac.uk



Short communication

A web-based tool for eliciting probability distributions from experts

David E. Morris^{a,*}, Jeremy E. Oakley^b, John A. Crowe^a^aElectrical Systems and Optics Division, Faculty of Engineering, University of Nottingham, Nottingham NG7 2RD, UK^bSchool of Mathematics and Statistics, University of Sheffield, Sheffield S3 7RH, UK

ARTICLE INFO

Article history:

Received 4 July 2013

Received in revised form

4 September 2013

Accepted 13 October 2013

Available online 12 November 2013

Keywords:

Bayesian prior distribution

Expert judgement

Subjective probability

Web-based elicitation

ABSTRACT

We present a web-based probability distribution elicitation tool: The MATCH Uncertainty Elicitation Tool. The Tool is designed to help elicit probability distributions about uncertain model parameters from experts, in situations where suitable data is either unavailable or sparse. The Tool is free to use, and offers five different techniques for eliciting univariate probability distributions. A key feature of the Tool is that users can log in from different sites and view and interact with the same graphical displays, so that expert elicitation sessions can be conducted remotely (in conjunction with tele- or videoconferencing). This will make probability elicitation easier in situations where it is difficult to interview experts in person. Even when conducting elicitation remotely, interviewers will be able to follow good elicitation practice, advise the experts, and provide instantaneous feedback and assistance.

© 2013 The Authors. Published by Elsevier Ltd. Open access under CC BY license.

Software availability

The MATCH Uncertainty Elicitation Tool: can be used for free and is available at <http://www.match.ac.uk/uncertainty/>.

1. Introduction

When modelling an environmental system, suitable data are not always available for obtaining robust estimates of all the model parameters. When data are unavailable or sparse, one option is to use expert judgement, and elicit probability distributions to represent uncertainty about each unknown parameter. These may form prior distributions that can be updated with data using Bayes' theorem, or the elicited distributions may be used directly with the model, if no data are available. An advantage of eliciting a probability distribution, rather than asking an expert for a single point estimate, is that it allows an expert to express his or her uncertainty about the parameter, and the consequences of this uncertainty can be investigated using sensitivity analysis techniques (Saltelli et al., 2008). O'Hagan (2012) gives an illustration of eliciting probability

distributions in modelling a carbon flux, and Krueger et al. (2012) give a broader discussion of the role of expert opinion in environmental modelling. Probability elicitation is of wide interest in many fields, and some textbooks on the subject include Morgan and Henrion (1990), Meyer and Booker (1991), Cooke (1991) and O'Hagan et al. (2006).

Eliciting a probability distribution is difficult, as the individual with subject matter expertise will typically not have experience in making probability judgements, and the process will be unfamiliar to him or her. O'Hagan et al. (2006) emphasize the important role of a facilitator: an individual with experience in probability elicitation, who is able to provide training and assistance to the expert. They also recommend giving feedback to the expert: highlighting any inconsistent judgements and, when fitting probability distributions to the expert's judgements, showing the fitted distribution to the expert to invite further reflection and check that the fitted distribution is acceptable. Typically, this requires the elicitation interview to be conducted face-to-face. However, in some cases, it may be difficult to interview the experts in person, or, when eliciting from multiple experts, to bring the experts together to the same location. In this paper, we present a software tool to overcome this difficulty (the software will also be useful in face-to-face elicitation).

In the next section, we briefly describe the elicitation process and the value of elicitation software. We then describe the MATCH Uncertainty Elicitation Tool, and how it enables users both to follow good practice, and conduct elicitation interviews remotely. The Tool

* Corresponding author. Tel.: +44 115 8467453.

E-mail address: david.morris@nottingham.ac.uk (D.E. Morris).

is named after the UK Research Council funded MATCH (Multidisciplinary Assessment of Technology for Healthcare) project. The development of the Tool was motivated by the need to elicit user opinions about medical devices, but we wish to promote the Tool across all modelling communities.

2. Eliciting a univariate probability distribution

We consider elicitation for a univariate uncertain quantity, which we denote by X . In most elicitation schemes, the expert is not asked to provide a probability distribution for X directly. Instead, the expert usually makes a small number of probability judgements about X , and the facilitator fits a probability distribution to these judgements. Typically, the expert specifies some points on the cumulative distribution function, either by providing some probabilities given some values of X , or by specifying some quantiles for X given some probabilities.

For example, in the quartile method, the expert provides a median $X_{0.5}$, lower quartile $X_{0.25}$, and upper quartile $X_{0.75}$. The facilitator can fit a probability distribution to these judgements using a least squares procedure. For example, if X is a proportion, a $\text{Beta}(a,b)$ distribution can be fitted by finding a and b to minimize $\{F(X_{0.25}; a, b) - 0.25\}^2 + \{F(X_{0.5}; a, b) - 0.5\}^2 + \{F(X_{0.75}; a, b) - 0.75\}^2$, where $F(X; a, b)$ is the cumulative distribution function of the $\text{Beta}(a, b)$ distribution. Once a and b have been chosen, the facilitator can provide feedback, by plotting the $\text{Beta}(a, b)$ density function, and by reporting additional quantiles from the $\text{Beta}(a, b)$ distribution, for example the 0.05 quantile and the 0.95 quantile.

The expert can be asked whether the chosen fit is acceptable, or whether he or she wishes to revise any judgements. For example, the expert may feel that the interval between the 0.05 and 0.95 quantiles is a little narrow, and that the initial judgements need revising. Ideally, the feedback and fitting is done in real time, so that the expert can either approve the chosen distribution or make revisions as necessary. Obviously, this requires software, and it is helpful for the facilitator to be able to discuss the fitted distribution with the expert.

3. The MATCH Uncertainty Elicitation Tool

The MATCH Tool is based on the SHELF elicitation package of Oakley and O'Hagan (2010). The SHELF package includes code written in R (R Development Core Team, 2013) for eliciting and fitting probability distributions. The Tool provides a web-based interface for the SHELF R code, which is more user-friendly, offers a little more flexibility for the elicitation methods, and includes a new feature for conducting elicitation remotely, which we discuss in Section 3.4. The user can specify probability judgements using the web interface, which links to the R code on a server, so that the users do not need to run R themselves. All outputs from the elicitation can be saved as a text file on the user's computer.

Similar web-based elicitation software, The Elicitor, has been presented in Bastin et al. (2013). The Elicitor offers more functionality for documenting the process, whereas the focus of the Tool is more on the elicitation itself, in that more elicitation methods are offered, as well as the facility for remote elicitation, which we think is the main highlight of the Tool.

3.1. Elicitation methods available in the Tool

There are five elicitation methods available, which we describe shortly. We believe them all to follow good practice as described in O'Hagan et al. (2006), in that they do not ask for difficult judgements such as means and variances, or for quantiles in the tails of the distribution such as the 0.05th and 0.95th, which can be hard to

judge without exhibiting over-confidence (Alpert and Raiffa, 1982). Users can experiment with the different methods to see which they prefer. In our experience, users tend to find the roulette method the easiest to use.

For plotting purposes, the expert must first specify lower and upper limits for the uncertain quantity X . We will refer to the interval specified by these limits as the range for X .

3.1.1. The roulette method

This method was suggested by Gore (1987) and reviewed in Johnson et al. (2010). The expert is provided with a grid of m equally sized bins which cover the range of X , and is asked to allocate a total of n chips between the bins. The probability of X lying in a particular bin is interpreted as the proportion of chips allocated to that bin. The Tool allows the user to specify any number of bins and use any number of chips. A screenshot of this method being used is given in Fig. 1.

3.1.2. The quartile method

This method is also referred to as the bisection method, and is illustrated in Raiffa (1968). The expert specifies the median, lower quartile and upper quartile of X . To help visualize these judgements, the expert can manipulate sliding bars. For example, to specify the median, the expert divides a coloured bar representing the range of X into two regions of equal probability.

3.1.3. The tertile method

The expert specifies the median, the 33rd percentile and the 66th percentile. The median is specified using a sliding bar as in the quartile method. The tertiles are specified by dividing a coloured bar into three equally likely sections.

3.1.4. The probability method

The expert specifies three probabilities. If the range of X is $[0, 1]$, then the default probabilities asked for are $P(0 < X < 0.25)$, $P(0.75 < X < 1)$ and $P(0 < X < 0.5)$, but the numbers in these inequalities can be changed. The probabilities can be specified by manipulating sliding bars, with the probability given by the proportion of the bar that is coloured.

3.1.5. The hybrid method

The expert specifies the median and two probabilities, for example $P(0 < X < 1/3)$ and $P(2/3 < X < 1)$. Again, these judgements can be specified by manipulating sliding bars.

3.2. Fitting and feedback

Once the judgements have been specified, the Tool fits various parametric distributions numerically, using the least squares procedure described above. The distributions available are normal, Student-t, beta (scaled if the parameter limits are not 0 and 1), gamma, log normal, and log Student-t. For both t-distributions, the degrees of freedom parameter is fixed at 3, to give a heavy-tailed alternative to the normal. The Tool will plot the fitted densities, and report which distribution fits the judgements the best. The Tool will also plot any two percentiles from the fitted distribution, which can be used as feedback. For example, if the roulette method has been used, the facilitator can then report the tertiles, and invite the expert to consider whether the three displayed intervals are equally likely. The results of all the distribution fitting can be exported as a text file.

3.3. The share option

Clicking the Share button on the Tool provides a weblink, and anyone viewing this weblink can then view and interact with the




Fig. 1. Eliciting a probability distribution using the roulette method. In the top graph, the expert expresses probability judgements by allocating chips to bins, and the bottom graph shows a Beta distribution which has been fitted to the expert's judgements.

same graphics. If the facilitator cannot interview the expert(s) in person, he or she can arrange a tele- or videoconference, share the weblink, and then the elicitation can be conducted remotely, with the facilitator still able to interact and advise the expert(s) on making and interpreting the probability judgements. We think this is the most attractive feature of the Tool; in particular, it makes eliciting from overseas experts more feasible.

For multiple experts, the Tool could be used to support a behavioural aggregation approach, in which the experts are asked to agree on a single set of probability judgements to represent the opinions of the group. In this situation, although any participant could interact with the displays, we would recommend that only the facilitator does so. The experts would still have the benefit of seeing the group's agreed probability judgements visualised and a distribution fitted, and any expert could request feedback as described in Section 3.3. The advantage in using the Tool would be that the group discussion, distribution fitting and feedback could all take place in a single session, in real time, with the facilitator advising the experts in probability assessment as necessary.

The Tool does not currently have facilities for eliciting distributions separately from each expert and calculating a pooled distribution. Of course, if the facilitator wishes to use mathematical aggregation, he or she could use the Tool to elicit a distribution from each expert separately, and then combine the results manually.

3.4. Data flow

The data flow of the Tool is described in Fig. 2. There is no limit to the number of web browsers used in an elicitation session at any one time, whilst the facilitator and experts can share the same browser for face-to-face eliciting. During the elicitation process,

inputs, which may be elicited judgements or other changes, made on any one browser are sent from that browser to the Web Server. Any elicited judgements are sent to the R code for fitting and the result is both stored in the database and sent back to the web browser. The input data itself is time stamped and stored in the database for use in playback mode (discussed in the next section) so that a history of how the elicited distribution was constructed can be recalled. Concurrently, the inputs are also sent to other browsers taking part in the elicitation session and these independently request the fitted distributions.

3.5. Playback mode

This is an experimental feature that we have added, that users may wish to use for reviewing an elicitation session. In playback mode, previous inputs are retrieved from the database and




Fig. 2. Data flow in the Tool. The web server enables a facilitator and expert to conduct an elicitation remotely, interacting with the same graphical displays. Fitting distributions to the expert's judgements is implemented using R.

transmitted to the playback browser, in the order that the inputs were received when eliciting took place. The inputs are transmitted in real time, that is with the same time intervals between successive inputs as the original elicitation (with a minimum time resolution of 1 s). This real time feature allows an audio recording of the elicitation session to be synchronised to the output on the screen when replaying. The browser behaves as though it is receiving data that has been inputted by another expert or facilitator, and displays them as if it is an inactive party present in the actual elicitation session. It is currently not indicated whether it was the expert or the facilitator who made each individual input during playback mode.

4. Discussion

We believe the Tool is a significant advance in making expert elicitation easier and more practical, particularly in cases when it is difficult to bring the experts and facilitator together to the same location. Of course, there is more to elicitation than simply having the right software; the role of the facilitator in training and assisting the experts is important, as is documenting the process carefully to provide transparency. Here, we recommend following the guidance set out in the templates provided within the SHELF package.

We hope analysts in various disciplines will find the tool helpful, and we are keen to get feedback and develop the tool accordingly.

Acknowledgements

The authors acknowledge support of this work through the MATCH Programme (EPSRC Grant EP/G012393/1), although the views expressed are entirely their own.

References

- Alpert, M., Raiffa, H., 1982. A progress report on the training of probability assessors. In: Kahneman, D., Slovic, P., Tversky, A. (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, UK, pp. 294–305.
- Bastin, L., Cornford, D., Jones, R., Heuvelink, G.B.M., Pebesma, E., Stasch, C., Nativi, S., Mazzetti, P., Williams, M., 2013. Managing uncertainty in integrated environmental modelling: the UncertWeb framework. *Environ. Model. Softw.* 39, 116–134.
- Cooke, R.M., 1991. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press, New York.
- Gore, S., 1987. Biostatistics and the Medical Research Council. *Med. Res. Council. News* 35, 19–20.
- Johnson, S.R., Tomlinson, G.A., Hawker, G.A., Granton, J.T., Grosbein, H.A., Feldman, B.M., 2010. A valid and reliable belief elicitation method for Bayesian priors. *J. Clin. Epidemiol.* 63, 370–383.
- Krueger, T., Page, T., Hubacek, K., Smith, L., Hiscock, K., 2012. The role of expert opinion in environmental modelling. *Environ. Model. Softw.* 36, 4–18.
- Meyer, M.A., Booker, J.M., 1991. *Eliciting and Analyzing Expert Judgment: a Practical Guide*. Academic Press, London.
- Morgan, M.G., Henrion, M., 1990. *Uncertainty: a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge.
- Oakley, J.E., O'Hagan, A., 2010. SHELF: the Sheffield Elicitation Framework (Version 2.0). School of Mathematics and Statistics, University of Sheffield. www.tonyohagan.co.uk/shelf.
- O'Hagan, A., 2012. Probabilistic uncertainty specification: overview, elaboration techniques and their application to a mechanistic model of carbon flux. *Environ. Model. Softw.* 36, 35–48.
- O'Hagan, A., Buck, C.E., Daneshkhah, A., Eiser, J.R., Garthwaite, P.H., Jenkinson, D., Oakley, J.E., Rakow, T., 2006. *Uncertain Judgements: Eliciting Experts' Probabilities*. Wiley, Chichester.
- R Development Core Team, 2013. *R: a Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- Raiffa, H., 1968. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Addison-Wesley.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S., 2008. *Global Sensitivity Analysis: the Primer*. Wiley, New York.