NPL
National Physical Laboratory

NPL REPORT MS 13

Software for GUM Supplement 2: User Manual

P M Harris and I M Smith

OCTOBER 2011

National
Measurement
System

NPL Report MS 13

Software for GUM Supplement 2: User Manual

P M Harris and I M Smith
Mathematics and Scientific Computing Group

October 2011

ABSTRACT

This report constitutes a *user manual* for software developed at the National Physical Laboratory to support the use of the second supporting document to the 'Guide to the expression of uncertainty in measurement' (GUM), GUM Supplement 2, concerned with methods implementing the propagation of distributions for measurement models with a general number of output quantities. The software enables users to apply the approaches to uncertainty evaluation described in GUM Supplement 2 to two of the example problems considered in the document. The software is intended to allow users to reproduce the results presented in tables and figures contained within GUM Supplement 2. It is also intended to help users learn about the methods for uncertainty evaluation described in the document by enabling them to experiment with (a) different information about the input quantities in the measurement models defining the example problems, and (b) different values for the parameters controlling the application of those methods.

NPL Report MS 13

ISSN 1754–2960

National Physical Laboratory,
Hampton Road, Teddington, Middlesex, United Kingdom TW11 0LW

We gratefully acknowledge the financial support of the UK Department for Business, Innovation and Skills (National Measurement Office)

Approved on behalf of NPLML by Dr Bajram Zeqiri,
Knowledge Leader for the Acoustics and Ionizing Radiation team

Contents

1 Introduction

1.1 Background

The 'Guide to the expression of uncertainty in measurement' (GUM) [2] is the primary document regarding the evaluation and reporting of uncertainty in measurement. Working Group 1, 'Expression of uncertainty in measurement', of the Joint Committee for Guides in Metrology (JCGM) is undertaking work to promote and extend the application of the GUM through the preparation of supporting documents [1]. The first such document, GUM Supplement 1 [3], is concerned with the use of a Monte Carlo method as an implementation of the propagation of distributions for uncertainty evaluation. A further document (GUM Supplement 2 [4]) is concerned with extensions of the GUM and GUM Supplement 1 to general measurement models, including models with a general number of output quantities.

GUM Supplement 2 includes four examples as follows:

1. Additive measurement model [4, clause 9.2];
2. Co-ordinate system transformation [4, clause 9.3];
3. Simultaneous measurement of resistance and reactance [4, clause 9.4];
4. Measurement of Celsius temperature using a resistance thermometer [4, clause 9.5].

The first two are examples of, respectively, linear and non-linear measurement models, each of which is bivariate, i.e., having two output quantities. The third example illustrates the treatment of a series of simultaneous indication values of a vector input quantity that have been obtained independently. The fourth example illustrates the treatment of implicit (univariate and multivariate) measurement models.

This report is a *user manual* for software developed at the National Physical Laboratory to support the use of GUM Supplement 2 by enabling users to apply the approaches to uncertainty evaluation described in this document to the first two of the above example problems. The software is intended to allow users to reproduce the results presented in tables and figures contained within GUM Supplement 2. The software is also intended to help users learn about the methods for uncertainty evaluation described in GUM Supplement 2. It enables them to experiment with (a) different information about the input quantities in the measurement models defining the example problems, e.g., estimates and associated standard uncertainties for the input quantities in those models, and (b) different values for the parameters controlling the application of those methods, e.g., the number of trials in an application of a Monte Carlo method.

The software does not allow users to undertake uncertainty calculations for measurement models other than those for the example problems considered. However, detailed information about the algorithms implemented in the software is available [5].

For each example problem, the software is used to apply the GUM uncertainty framework [4, clause 6] and a Monte Carlo method [4, clause 7] as approaches to the propagation and summarizing stages of uncertainty evaluation [4, subclause 5.2]. Each approach delivers

- an estimate $\boldsymbol{y} = (y_1, y_2)^\top$ of the output quantity $\boldsymbol{Y} = (Y_1, Y_2)^\top$ in the bivariate measurement model,
- the standard uncertainties $u(y_1)$ and $u(y_2)$ and correlation $r(y_1, y_2)$ associated with the estimate,
- coverage factors defining elliptical and rectangular coverage regions for $\boldsymbol{Y}$ corresponding to a specified coverage probability p, and
- approximations to the marginal probability density functions (PDFs) for Y_1 and Y_2 and the joint PDF for $\boldsymbol{Y}$.

The Monte Carlo procedure may be applied either non-adaptively or adaptively. In a non-adaptive application of the procedure [4, subclause 7.2], a fixed number of Monte Carlo trials, specified by the user, is undertaken and tests of whether the results obtained have stabilized in a statistical sense performed at the end of the calculation. The tests are based on numerical tolerances for the results calculated in terms of a number $n_{\rm stab}$ of significant decimal digits regarded as meaningful in the values of the results [4, subclause 7.8.2].[1] In an adaptive application of the procedure [4, subclause 7.8], the calculation is terminated when either (a) the results obtained have stabilized, or (b) a maximum number of trials, specified by the user, has been undertaken.

In the case that the results obtained from the Monte Carlo procedure have stabilized, those results are used to validate the results obtained from the GUM uncertainty framework [4, clause 8]. The basis of the validation are numerical tolerances for the results calculated in terms of a number $n_{\rm val}$ of significant decimal digits regarded as meaningful in the values of the results [4, subclause 7.8.2].[2] A sufficient number of Monte Carlo trials should be undertaken in obtaining results from a Monte Carlo method for the purpose of validating those from the GUM uncertainty framework, and it is recommended that the numerical tolerances set in the adaptive procedure are no greater than one fifth of the respective numerical tolerances set in the validation procedure.

Software to support the use of GUM Supplement 1 is also available, and is described in [6].

The report is organized as follows. Section 2 gives information on installing and uninstalling the software. Section 3 describes how to use the software. General information is

[1] The notation $n_{\rm stab}$ is used in place of the generic notation $n_{\rm dig}$ used in subclause 7.8.2 of GUM Supplement 2 in the context of a test of stabilization of the results obtained from a Monte Carlo method. See section 3.1.5.

[2] The notation $n_{\rm val}$ is used in place of the generic notation $n_{\rm dig}$ used in subclause 7.8.2 of GUM Supplement 2 in the context of validating the results obtained from the GUM uncertainty framework. See section 3.1.5.

provided relating to the software for both example problems (section 3.1), as well as information specific to each problem and its associated software (sections 3.2 and 3.3). Section 4 describes the conditions of use of the software that give the results presented in tables and figures contained within GUM Supplement 2.

1.2 Software User Licence Agreement

The software is provided with a Software User Licence Agreement (Ref: MSC/L/11/002) and the use of the software is subject to the terms laid out in that agreement. By installing and running the software, the user accepts the terms of the agreement.

To run the software, the user must install MATLAB's Component Runtime (MCR) libraries (section 2). The user must accept the terms of the MCR Library License as part of the installation of the MCR libraries.

1.3 Operating the software

The operation of the software, and in particular the numerical and graphical display of results, is intended for the example problems as they are described in GUM Supplement 2, and for instances of those problems that are 'close' to those described therein. Therefore, for each example problem, the values of the expectation and standard deviation of each input quantity in the model for the problem, obtained from the probability distribution that characterizes the quantity, are required to lie within stated intervals. The software undertakes checks on the values set by the user, and will not accept values lying outside those intervals (section 3.1.1). However, the performance of the software and the adequacy of the display of results are not guaranteed. For example, if the user sets a very large value for the number of Monte Carlo trials, the time for the computation to be completed may be unacceptably long. Also, if the user chooses a very small value for the standard uncertainty associated with the estimate of an input quantity, the number of figures to which the value is displayed may be inadequate to distinguish the displayed value from zero. Generally, the number of figures to which results are displayed is greater than would normally be used to report results, but the intention is to facilitate the comparison of results. Sections 3.2 and 3.3 give, for the two example problems, the default values and settings that define the problems, and the intervals in which the expectations and standard deviations for each input quantity are required to lie.

2 Installing and uninstalling the software

The software takes the form of two application programs, corresponding to the first two example problems listed in section 1.1, called

1. `NPLUnc_commoneffect.exe`,[3] and
2. `NPLUnc_coordinatetransform.exe`.

The application programs have been created by compiling (using the MATLAB compiler) software implemented in the MATLAB programming language [7]. The programs have been created and tested on a personal computer running the Microsoft Windows XP Professional operating system.

To run the application programs it is first necessary to install MATLAB's Component Runtime (MCR) libraries. This is done by running the MCR installation program

```
MCRInstaller.exe
```

once on the target machine, i.e., the machine on which it is intended to run the application programs. It is necessary to have administrative privileges for the target machine because both the system registry and system path are modified as part of the installation process. The MCR installation program installs the MCR libraries, registers the components as needed, and updates the system path to point to the MCR binary directory. The installation process takes some time due to the number of files that are installed. The MCR installation program is about 173 MB in size, and the installed libraries require about 456 MB of disk space.

The software is uninstalled by

- running the MCR installation program `MCRInstaller.exe` and selecting 'Remove' to uninstall the MCR libraries, and
- deleting the application programs.

3 Using the software

3.1 General

An application program is run either by (a) double-clicking on the corresponding executable file (with the extension `.exe`) in Windows Explorer, or (b) opening a DOS window, navigating to the folder containing the program, typing the name of the program (without the extension `.exe`), and pressing `Return`. Running an application program the first time can take longer than when the program is run subsequently.

Each program allows the user to perform the following generic operations:

[3] The name `NPLUnc_commoneffect.exe` is chosen to distinguish the program from the program `NPLUnc_additive.exe`, which is part of the software released to support GUM Supplement 1 [6].

1. Modify the PDFs used to characterize the input quantities in the measurement model;

2. Modify the coverage probability;

3. Run an uncertainty calculation;

4. Exit.

Each of these operations is described below. Information specific to each example problem and its associated program is presented in sections 3.2 and 3.3.

3.1.1 Modify the PDFs

For example problem 1 (additive measurement model), each input quantity $X_i,\ i = 1, 2, 3$, in the measurement model is characterized by a probability distribution, which may be:

1. A Gaussian distribution $\mathrm{N}(\mu, \sigma^2)$ with expectation μ and standard deviation $\sigma \geq 0$ [3, subclause 6.4.7];

2. A rectangular distribution $\mathrm{R}(a, b)$ with lower limit a and upper limit $b \geq a$ [3, subclause 6.4.2].

The input quantities are mutually independent.

For each input quantity, the user is presented with a list of (two) probability distributions with a distribution highlighted. The distribution highlighted is the 'current' distribution that characterizes the input quantity and corresponds to either a default distribution (if the PDF has not been modified previously by the user) or a distribution that has been set by the user. This distribution is defined by 'current' values of the parameters for the distribution which either take default values (if the PDF has not been modified previously by the user) or have been set by the user. Section 3.2 gives the default distributions and default parameter values for example problem 1. The current distribution with its current parameter values is selected by pressing `Cancel`.

The PDF used to characterize the input quantity is modified by selecting a distribution from the list and pressing `OK`. Values for the parameters defining the selected distribution are then set as follows:

- When the current distribution is selected, the user may change the current values for the parameters of that distribution. New values are set by changing one or more values and pressing `OK`. The current values are restored in the text boxes by pressing `Cancel`.

- When a distribution other than the current distribution is selected, the user must set values for the parameters of the distribution. Values are set by typing in the values and pressing `OK`. Values previously entered in the text boxes are cleared by pressing `Cancel`.

The following checks are undertaken to ensure the parameter values are valid:

1. $\sigma \geq 0$ for a Gaussian distribution;
2. $b \geq a$ for a rectangular distribution.

Furthermore, to ensure that the problem defined by the user is 'close' to the (default) problem described in GUM Supplement 2, checks are undertaken to ensure that the expectation and standard deviation of each input quantity lie within stated intervals. These checks are described in section 3.2. If a check on the probability distribution set by the user for an input quantity fails, a warning message is displayed and the user must set another distribution for the quantity.

Final checks are undertaken to ensure that the standard uncertainties associated with the estimates of the input quantities are such that the covariance matrix associated with the estimates of the output quantities is positive definite. This condition ensures that coverage regions for the output quantities, which depend on the inverse of the covariance matrix, are well defined. The checks require that:

1. Each output quantity depends on at least one input quantity for which the standard uncertainty associated with the estimate of the quantity is strictly positive, i.e., it must not hold that $u(x_1) = u(x_3) = 0$ or $u(x_2) = u(x_3) = 0$ or $u(x_1) = u(x_2) = u(x_3) = 0$;[4]
2. There is at least one input quantity for which the standard uncertainty associated with the estimate of the quantity is strictly positive *and* the quantity is not common to both output quantities, i.e., it must not hold that $u(x_1) = u(x_2) = 0$.[5]

If either of these final checks fails, a warning message is displayed and the user must set new distributions for all the input quantities.

For example problem 2 (co-ordinate system transformation), the (vector) input quantity $\boldsymbol{X} = (X_1, X_2)^{\mathrm{T}}$, with $\boldsymbol{x} = (x_1, x_2)^{\mathrm{T}}$ an estimate of $\boldsymbol{X}$ with associated covariance matrix $\boldsymbol{U_x}$, is characterized by a bivariate Gaussian distribution $\mathrm{N}(\boldsymbol{x}, \boldsymbol{U_x})$. The covariance matrix $\boldsymbol{U_x}$ is defined by

$$\boldsymbol{U_x} = \begin{bmatrix} u^2(x_1) & r(x_1, x_2)u(x_1)u(x_2) \\ r(x_1, x_2)u(x_1)u(x_2) & u^2(x_2) \end{bmatrix},$$

[4]This check ensures that for each output quantity the standard uncertainty associated with an estimate of the quantity is strictly positive.

[5]This check ensures that the correlation of the output quantities is not ± 1.

where $u(x_i)$ is the standard uncertainty associated with x_i, and $r(x_1, x_2)$ is the correlation coefficient associated with x_1 and x_2.

The user is presented with a set of default values for $x_1, u(x_1), x_2, u(x_2)$ and $r(x_1, x_2)$, and may change these values. New values are set by changing one or more values and pressing `OK`. The default values are restored in the text boxes by pressing `Cancel`. The following checks are undertaken to ensure the parameters of the distribution characterizing $\boldsymbol{X}$ are valid:

1. $u(x_i) > 0,\ i = 1, 2$;
2. $-1 < r(x_1, x_2) < 1$.

Finally, checks are undertaken to ensure the conditions given in section 3.3 hold. If a check fails, a warning message is displayed and the user must set different values for the parameters of the probability distribution characterizing the input quantities.

3.1.2 Modify the coverage probability

The user may set the coverage probability to be 0.90, 0.95 or 0.99. The user is presented with a list of values with the default value of 0.95 highlighted. The coverage probability is set by selecting a value from the list and pressing `OK`. The coverage probability is set to its default value, irrespective of whether a different value has been selected, by pressing `Cancel`.

3.1.3 Run an uncertainty calculation

The following steps are undertaken:

1. Information is displayed about the input quantities in the measurement model. The information is either default information set when the program is first run, or that set by the user as part of the operation ‘Modify the PDFs’ (section 3.1.1).
2. The GUM uncertainty framework for uncertainty evaluation is applied to obtain an estimate $\boldsymbol{y} = (y_1, y_2)^\top$ of the output quantity $\boldsymbol{Y} = (Y_1, Y_2)^\top$, the standard uncertainties $u(y_1)$ and $u(y_2)$ and correlation $r(y_1, y_2)$ associated with the estimate, and coverage factors defining elliptical and rectangular coverage regions for $\boldsymbol{Y}$. The coverage probability for the coverage region takes either a default value set when the program is first run, or a value set by the user as part of the operation ‘Modify the coverage probability’ (section 3.1.2).
3. The results obtained from an application of the GUM uncertainty framework are displayed.

4. The user may set various controls for an application of a Monte Carlo method for uncertainty evaluation. The controls include:

   (a) Whether the Monte Carlo procedure is applied adaptively (`y` or `Y`) or non-adaptively (`n` or `N`). If the value set by the user is not one of `y`, `Y`, `n` or `N`, a warning message is displayed and the user must set a different value.

   (b) The maximum number of trials undertaken in an application of the Monte Carlo procedure. The number of trials is specified as a multiple of 10^4. In a non-adaptive application of the procedure, the maximum number of trials will be undertaken. In an adaptive application, fewer trials will be undertaken if the results obtained from the procedure have stabilized before the maximum number of trials has been reached. The value set by the user is rounded to the nearest integer M_h, and M_h must satisfy $M_h \geq 10$; otherwise a warning message is displayed and the user must set a different value.

   (c) The number of bins (or interval classes) for displaying as a scaled frequency distribution the approximation to the (marginal) PDF for each output quantity obtained from the procedure. The value set by the user is rounded to the nearest integer N_{bin}, and N_{bin} must satisfy $N_{\text{bin}} \geq 10$; otherwise a warning message is displayed and the user must set a different value. The number of bins (or rectangular classes) used to display as a scaled frequency distribution the approximation to the (joint) PDF for the bivariate output quantity is $N_{\text{bin}} \times N_{\text{bin}}$.

   (d) The initial state of the pseudo-random number generator used by the Monte Carlo procedure. The value set by the user is rounded to the nearest integer S, and S must satisfy $0 \leq S \leq 2^{32} - 1$; otherwise a warning message is displayed and the user must set a different value.

   (e) The number n_{stab} of significant decimal digits used to determine numerical tolerances for testing the stabilization of the results obtained from the Monte Carlo procedure (section 3.1.5). The value set by the user must satisfy $n_{\text{stab}} \geq 1$; otherwise a warning message is displayed and the user must set a different value.

   (f) The number n_{val} of significant decimal digits used to determine numerical tolerances for testing whether the results obtained from an application of the GUM uncertainty framework are validated by those obtained from an application of the Monte Carlo procedure (section 3.1.5). The value set by the user must satisfy $n_{\text{val}} \geq 1$; otherwise a warning message is displayed and the user must set a different value.

   (g) The number of times contour smoothing is applied for displaying an approximation to the PDF for the output quantities (section 3.1.6). The value set by the user is rounded to the nearest integer n_{sm}, and n_{sm} must satisfy $n_{\text{sm}} \geq 0$; otherwise a warning message is displayed and the user must set a different value.

   Default values are provided for each of these controls, except the initial state of the pseudo-random number generator, for which a different value is set each time a calculation is run. New values are set by changing one or more values and pressing `OK`.

The controls are set to their default values and the calculation is run, irrespective of whether a value has been changed, by pressing `Cancel`.

5. A Monte Carlo method for uncertainty evaluation is applied to obtain an estimate $\boldsymbol{y}$ of $\boldsymbol{Y}$, the standard uncertainties $u(y_1)$ and $u(y_2)$ and correlation $r(y_1, y_2)$ associated with the estimate, and coverage factors defining elliptical and rectangular coverage regions for $\boldsymbol{Y}$. The coverage probability is as for the application of the GUM uncertainty framework (step 2). A progress bar gives an indication of the progress made by the Monte Carlo calculation. In an adaptive application of the Monte Carlo procedure that requires fewer than the maximum number of trials for the results to stabilize, the calculation will complete before the progress bar is full.

6. The results obtained from an application of a Monte Carlo method are displayed. In addition, the following information is displayed:

   (a) The value of the coverage probability set by the user (section 3.1.2).

   (b) Information about whether the Monte Carlo calculation has stabilized. The information includes the number $h \times 10^4$ $(h \leq M_h)$ of trials undertaken and the number $n_{\rm stab}$ of significant decimal digits used to determine numerical tolerances for testing the stabilization of the results. It also includes the values $2s_{y_1}$, $2s_{u(y_1)}$, $2s_{y_2}$, $2s_{u(y_2)}$, $2s_{\lambda_{\max}}$, $2s_{k_p}$ and $2s_{k_q}$ that are compared with, respectively, numerical tolerances δ_1, δ_1, δ_2, δ_2, ρ, κ_p and κ_q in the test [4, subclause 7.8.3]. Here, the value s_{y_1} is the standard deviation associated with the average of the estimates $y_1^{(1)}, y_1^{(2)}, \ldots, y_1^{(h)}$ of the output quantity Y_1 obtained from h sets of 10^4 trials, and similarly for $s_{u(y_1)}$, s_{y_2}, $s_{u(y_2)}$, $s_{\lambda_{\max}}$, s_{k_p} and s_{k_q}. The computation is regarded as having stabilized when *all* the values $2s_{y_1}$, $2s_{u(y_1)}$, $2s_{y_2}$, $2s_{u(y_2)}$, $2s_{\lambda_{\max}}$, $2s_{k_p}$ and $2s_{k_q}$ are no greater than the corresponding tolerances. In the case that the computation has not stabilized, an examination of the values can help to decide how many more trials might be needed to achieve stabilization of the results for the chosen value of $n_{\rm stab}$.

   (c) In the case that the Monte Carlo calculation has stabilized, information about whether the results obtained from an application of the GUM uncertainty framework are validated by those from an application of the Monte Carlo procedure, and the number $n_{\rm val}$ of significant decimal digits used to determine numerical tolerances for testing the validation of the results.

7. For each component of $\boldsymbol{Y}$, a graph is provided showing the marginal PDFs obtained from the applications of the GUM uncertainty framework and a Monte Carlo method. The marginal PDF provided by the GUM uncertainty framework is shown as a Gaussian distribution, and that provided by a Monte Carlo method as a scaled frequency distribution based on $N_{\rm bin}$ bins (or interval classes).

8. Graphs are provided showing the joint PDF for $\boldsymbol{Y}$, and corresponding elliptical and rectangular coverage regions, obtained from the application of the GUM uncertainty

framework. The joint PDF is displayed as either a coloured contour plot[6] or a coloured checkerboard plot[7] determined from values of the PDF, which takes the form of a bivariate Gaussian distribution, for values of $\boldsymbol{Y}$ defining a uniform grid on the rectangular region $[y_1 - 3u(y_1), y_1 + 3u(y_1)] \times [y_2 - 3u(y_2), y_2 + 3u(y_2)]$. The boundaries of the elliptical and rectangular coverage regions are displayed in blue and red, respectively.

9. Graphs are provided showing (the approximation to) the joint PDF for $\boldsymbol{Y}$, and corresponding elliptical, rectangular and smallest coverage regions, obtained from the application of a Monte Carlo method. The joint PDF is displayed as either a coloured contour plot or a coloured checkerboard plot determined from values of the PDF, which are approximated by scaled frequencies values, for $N_{\rm bin} \times N_{\rm bin}$ bins (or rectangular classes). The boundaries of the elliptical and rectangular coverage regions are displayed in blue and red, respectively. The smallest coverage region, which is composed of a subset of the bins used to define the joint PDF,[8] is displayed as a 'grey filled-in' region.

10. The user may control the way the joint PDFs are displayed by choosing from the following options:

    (a) *Coloured contour plots with default colour scales.* For each PDF eight contours are displayed corresponding to probability density values that are uniformly spaced between zero and the maximum probability density value for the PDF (excluding probability density values equal to zero and the maximum probability density value). Contours with the same colour in the two graphs are generally associated with different probability density values. This option is the default choice for displaying the joint PDFs.

    (b) *Coloured contour plots with the same colour scale.* For each PDF contours are determined corresponding to eight probability density values that are uniformly spaced between zero and the maximum probability density value calculated for both PDFs (excluding probability density values equal to zero and the maximum probability density value). For the PDF for which the maximum probability density value is attained, all eight contours will be displayed. The other PDF may be displayed using less than eight contours and, if it is appreciably different, using considerably less than eight contours. Contours with the same colour in the two graphs are associated with the same probability density value. This option allows a comparison of the two PDFs to be made.

    (c) *Coloured checkerboard plots with default colour scales.* For each PDF the same colour map (or set of colours) is used, but the same colour in the two graphs is generally associated with different probability density values.

[6]In a coloured contour plot each contour is defined by a set of values of $\boldsymbol{Y}$ for which the joint PDF takes the same probability density value.

[7]In a coloured checkerboard plot values of $\boldsymbol{Y}$ for which the joint PDF takes the same probability density value are associated with the same 'colour' from a set of colours that define a 'colour map'.

[8]The 'resolution' of the smallest coverage region is determined by the number $N_{\rm bin}$ of bins chosen.

(d) *Coloured checkerboard plots with the same colour scale.* For each PDF the same colour map is used, and the same colour in the two graphs is associated with the same probability density value. This option also allows a comparison of the two PDFs to be made.

For each displayed graph, a 'colour bar' may be included by pressing the 'insert colorbar' (toggle) button on the figure toolbar. The colour bar provides a scale for probability density in the graph.

3.1.4 Exit

Pressing `Exit` will terminate the program and close all windows associated with running the program.

3.1.5 Setting the numerical tolerances

Let $n_{\rm dig}$ denote the number of significant decimal digits regarded as meaningful in a numerical value z. $n_{\rm dig} = n_{\rm stab}$ when testing for the stabilization of the results from a Monte Carlo calculation, and $n_{\rm dig} = n_{\rm val}$ when validating the results obtained from an application of the GUM uncertainty framework against those obtained from an application of the Monte Carlo procedure. The numerical tolerance δ associated with z is defined as follows [4, subclause 7.8.2.1]:

1. Express z in the form $c \times 10^{\ell}$, where c is an $n_{\rm dig}$ decimal digit integer and ℓ is an integer;
2. Form
$$\delta = \frac{1}{2} 10^{\ell}. \tag{1}$$

For example, consider the case that $z = 1$. If two significant decimal digits are regarded as meaningful in z, then $n_{\rm dig} = 2$, $z = 10 \times 10^{-1}$ from which $\ell = -1$ and $\delta = 0.05$. Similarly, if three significant decimal digits are regarded as meaningful, then $n_{\rm dig} = 3$, $z = 100 \times 10^{-2}$ from which $\ell = -2$ and $\delta = 0.005$.

The software uses the following formula to evaluate ℓ:

$$\ell = {\rm floor}(\log_{10} z) - (n_{\rm dig} - 1), \tag{2}$$

in which ${\rm floor}(z)$ rounds the value z to the nearest integer towards $-\infty$. Thus, in the above example, with $z = 1$ and $n_{\rm dig} = 2$, $\ell = {\rm floor}(\log_{10} 1) - (2 - 1) = -1$, and when $n_{\rm dig} = 3$, $\ell = {\rm floor}(\log_{10} 1) - (3 - 1) = -2$, as required. However, the formula implemented in the software is not restricted to values of $n_{\rm dig}$ and ℓ that are integers. For example, setting $n_{\rm dig} = 2.7$ gives $\ell = -1.7$ and $\delta \approx 0.01$.

It follows from formulae (1) and (2) that

$$n_{\mathrm{dig}} = \mathrm{floor}(\log_{10} z) - \log_{10} 2\delta + 1.$$

This formula can be used to set n_{dig}, which is required by the software, in terms of δ.

For the different results obtained from an application of the GUM uncertainty framework and a Monte Carlo method, the corresponding numerical tolerances are set in the following way:

- The numerical tolerance corresponding to an estimate y_j and the standard uncertainty $u(y_j)$ associated with y_j is calculated in terms of a number of significant decimal digits regarded as meaningful in the value of $u(y_j)$ [4, subclause 7.8.2.2];
- The numerical tolerance corresponding to the correlation coefficients $r(y_i, y_j)$ associated with the estimates y_i and y_j is calculated in terms of a number of significant decimal digits in the value of the largest eigenvalue of the matrix of correlation coefficients [4, subclause 7.8.2.3];
- The numerical tolerance corresponding to a coverage factor k_p (or k_q) used to specify an elliptical (or rectangular) coverage region is calculated in terms of a number of significant decimal digits regarded as meaningful in the value of k_p (or k_q) [4, subclause 7.8.2.7].

3.1.6 Contour smoothing

When displaying the joint PDF obtained from a Monte Carlo method, the user may choose first to smooth the scaled frequency values that define an approximation to that PDF. As a consequence a smoother visualization, e.g., smoothed contours in a contour plot of the joint PDF, are obtained. A smoothed value $\widetilde{p}_{ij}$ of a scaled frequency value p_{ij} is evaluated as a weighted average of p_{ij} and its neighbouring values as follows:

$$\widetilde{p}_{ij} = \sum_{k=-1}^{1} \sum_{\ell=-1}^{1} w_{k\ell} p_{i+k,j+\ell},$$

where

$$\begin{aligned} w_{-1,-1} = w_{-1,1} = w_{1,-1} = w_{1,1} &= 0.05, \\ w_{-1,0} = w_{0,-1} = w_{0,1} = w_{1,0} &= 0.10, \\ w_{0,0} &= 0.40. \end{aligned}$$

When a value $p_{i+k,j+\ell}$ does not exist, e.g., when the index $i + k < 1$, the value is replaced by zero. To control the degree of smoothing, the user can set the number n_{sm} of times smoothing is applied.

3.2 Additive measurement model

The additive measurement model takes the form [4, subclause 9.2.1]

$$Y_1 = X_1 + X_3, \quad Y_2 = X_2 + X_3.$$

The default settings for the probability distributions for the input quantities are that each input quantity $X_i,\ i = 1, 2, 3$, is characterized by the standard Gaussian distribution $\mathrm{N}(0, 1)$. In the case that the probability distributions are modified by the user, the estimates (expectations) x_i of X_i and associated standard uncertainties (standard deviations) $u(x_i)$ are required to satisfy:

- $x_i \in [-10, 10],\ i = 1, 2, 3$;
- $u(x_i) \leq 10,\ i = 1, 2, 3$.

The default value for the coverage probability is $p = 0.95$.

The default settings for the application of the Monte Carlo procedure are:

- The application of the procedure is adaptive;
- $M_h = 1\,000$;
- $N_{\mathrm{bin}} = 75$;
- $n_{\mathrm{stab}} = 3$;
- $n_{\mathrm{val}} = 2$;
- $n_{\mathrm{sm}} = 0$.

3.3 Co-ordinate system transformation

The measurement model takes the form [4, subclause 9.3.1.1]

$$Y_1^2 = X_1^2 + X_2^2, \quad \tan Y_2 = X_2/X_1.$$

The default setting for the (joint) probability distribution for the input quantities is that the vector input quantity $\boldsymbol{X} = (X_1, X_2)^{\mathrm{T}}$ is characterized by a bivariate Gaussian distribution $\mathrm{N}(\boldsymbol{x}, \boldsymbol{U_x})$ with $x_1 = 0.10$, $x_2 = 0.00$, $u(x_1) = u(x_2) = 0.01$ and $r(x_1, x_2) = 0$. In the case that the probability distribution is modified by the user, the estimates (expectations) x_i of X_i and associated standard uncertainties (standard deviations) $u(x_i)$ are required to satisfy:

- $x_i \in [-0.5, 0.5],\ i = 1, 2$, and x_1 and x_2 cannot both be zero;

- $u(x_i) \leq 0.1,\ i = 1, 2.$

The default value for the coverage probability is $p = 0.95$.

The default settings for the application of the Monte Carlo procedure are:

- The application of the procedure is non-adaptive;
- $M_h = 1\,000$;
- $N_{\text{bin}} = 100$;
- $n_{\text{stab}} = 3$;
- $n_{\text{val}} = 2$;
- $n_{\text{sm}} = 2$.

4 Reproducing the results given in GUM Supplement 2

To reproduce the numerical and graphical results for the first two example problems given in GUM Supplement 2 it is necessary to set appropriately (a) the distributions for the input quantities, (b) the required coverage probability, and (c) the controls for the application of a Monte Carlo method including (d) the value S of the initial state of the pseudo-random number generator used by the software. Regarding (a), (b) and (c), the default values and settings listed in sections 3.2 and 3.3 provide a starting point. Regarding (d), tables 1 and 2 give the values S that must be set by the user. For example:

1. For the additive measurement model [4, subclause 9.2], to obtain the numerical results contained in tables 3, 4 and 5 of GUM Supplement 2, use the values for S given in the corresponding rows of table 1. For example, to obtain the results in rows 5 and 6 (Adaptive MCM) of table 3 in GUM Supplement 2, use the default values listed in section 3.2 and $S = 2\,108\,072$ (row 5) or $S = 2\,080\,974$ (row 6).

2. For the problem of co-ordinate system transformation [4, subclause 9.3], to obtain the numerical results in tables 6 and 7 of GUM Supplement 2, use the values for S given in table 2. For example, to obtain the results for the case $r(x_1, x_2) = 0$ and $x_1 = 0.10$ given in the last two rows of table 6 in GUM Supplement 2, use the default values listed in section 3.3 and $S = 2\,119\,134$.

Method	Initial state S
MCM $(M = 10^5)$	2 105 312
MCM $(M = 10^6)$	2 069 108
MCM $(M = 10^7)$	2 119 134
Adaptive MCM	2 108 072
Adaptive MCM	2 080 974

Table 1: Initial states of the pseudo-random number generator used by the Monte Carlo procedure for the additive measurement model [4, subclause 9.2].

Method	$r(x_1, x_2)$	x_1	Initial state S
MCM $(M = 10^7)$	0.0	0.001	2 104 850
MCM $(M = 10^7)$	0.0	0.010	2 064 567
MCM $(M = 10^7)$	0.0	0.100	2 093 964
MCM $(M = 10^7)$	0.9	0.001	2 129 001
MCM $(M = 10^7)$	0.9	0.010	2 107 709
MCM $(M = 10^7)$	0.9	0.100	2 097 554

Table 2: Initial states of the pseudo-random number generator used by the Monte Carlo procedure for the example problem of co-ordinate system transformation [4, subclause 9.4].

Acknowledgements

The work described here was supported by the National Measurement Office of the UK Department of Business, Innovation and Skills as part of its NMS Mathematics and Modelling for Metrology programme.

References

[1] Bich W, Cox M G and Harris P M 2006 Evolution of the 'Guide to the expression of uncertainty in measurement' *Metrologia* **43** S161–7

[2] BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP, and OIML *Evaluation of measurement data — Guide to the expression of uncertainty in measurement* Joint Committee for Guides in Metrology, JCGM 100:2008 `http://www.bipm.org/utils/common/documents/jcgm/JCGM_100_2008_E.pdf`

[3] BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP and OIML *Evaluation of measurement data — Supplement 1 to the 'Guide to the expression of uncertainty in measurement' — Propagation of distributions using a Monte Carlo method* Joint Committee for Guides in Metrology, JCGM 101:2008 `http://www.bipm.org/utils/common/documents/jcgm/JCGM_101_2008_E.pdf`

[4] BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP and OIML *Evaluation of measurement data — Supplement 2 to the 'Guide to the expression of uncertainty in measurement' — Extension to any number of output quantities* Joint Committee for Guides in Metrology, JCGM 102:2008

[5] Cox M G, Harris P M and Smith I M 2010 Software specification for uncertainty evaluation *NPL report MS 7*, National Physical Laboratory, Teddington, UK

[6] Cox M G, Harris P M and Smith I M 2008 Software for GUM Supplement 1: User manual *NPL report MS 1*, National Physical Laboratory, Teddington, UK

[7] MATLAB (Copyright The MathWorks, Inc.) `www.mathworks.com`